UCL

In Touch with the Wild: Exploring Real-time Feedback for Learning to Play the Violin

Rose Mary Grace Johnson

Thesis submitted in partial fulfilment of the

requirements for the degree of

Doctor of Philosophy

of

University College London

Department of Computer Science

University College London

2014

Declaration

I, Rose Mary Grace Johnson, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Abstract

Real-time feedback has great potential for enhancing learning complex motor-skills by enabling people to correct their mistakes as they go. Multimodal real-time cues could provide reinforcement to inform players whether they are making the correct or incorrect movements at a given time. However, little is known about how best to communicate information in real-time so that people can readily perceive and apply it to improving their movement while learning complex motor-skills. This thesis addresses this gap in knowledge by investigating how real-time feedback can enhance learning to play the violin. It explores how haptic and visual feedback are perceived, understood and acted upon in real-time when engaged in the primary task of playing the violin.

Prototypes were built with sensors to measure movement and either vibrations on the body or visual signals as feedback. Three in-the-wild user studies were conducted: one comparing visual and vibrotactile feedback for individual practice; one investigating shared feedback at a musical summer school; and one examining real-time feedback as part of a programme of learning at a high school. In-the-wild studies investigate users interacting with technology in a naturalistic setting, with all the demands that this entails. The findings show real-time feedback is effective at improving violin technique and can support learning in other ways such as encouraging mutual support between learners. The positive learning outcomes, however, need to be understood with respect to the complex interplay between the technology, demands of the setting and characteristics of individual learners. A conceptual framework is provided that outlines these interdependent factors. The findings are discussed regarding their applicability to learning other physical skills and the challenges and insights of using an in-the-wild methodology. The contribution of this thesis is to demonstrate empirically and theoretically how real-time vibrotactile and visual feedback can enhance learning a complex motor-skill.

Acknowledgements

I would like to thank my PhD supervisors, Yvonne Rogers, Janet van der Linden and Nadia Bianchi-Berthouze for always being positive and supportive. This has been an incredible experience which you have guided me through and taught me new ways of thinking. Thank you for always encouraging me to go a step further than I thought I was able to go.

Thank you to the Open University for giving me the opportunity to study such an interesting topic and for funding the first two years of my PhD. Thank you to UCL for allowing me to transfer my PhD when Yvonne moved to UCLIC and for funding the remainder of my PhD. Special thanks to Melanie Johnson, Dawn Bailey and JJ Giwa for making sure I had support through my maternity leave and ensuring that I could return part-time.

Thank you to the organisers of the musical summer school and the teachers at the high school for helping me to conduct my research. Thank you to all the people who participated in the studies and the experts and teachers who have advised me on the design. Thank you all for giving up your time to try these strange prototypes for practising the violin. Thank you for believing in them, and improving them, and most of all for sharing your love of music with me. I have only avoided naming you all individually to protect anonymity.

I'd like to thank Jon Bird and Erwin Schoonderwaldt for giving me an amazing start to my PhD with the MusicJacket project. Thank you to Paul Marshall for teaching me statistics with SPSS. Thanks to all my friends and colleagues at the Open University and UCL for all the chats over lunch that have stimulated many ideas that have gone into this PhD. Thanks to Clara Mancini and Anna Cox for your support - especially Anna who gave me helpful advice when I came back from maternity leave.

Thank you to everyone who helped to proof-read my thesis: Lorna Wall, my brother Tom, my Dad and Lorraine.

Thank you to everyone who has helped me with Dylan while I have been working on my PhD. Thank you to my Mum and Dad for looking after Dylan so regularly. Thank you to my dear friend Tina and little Karali for giving Dylan wonderful days playing at your house while I have been away at UCL. Thank you to Hannah and Logan for taking Dylan for sunny days in the garden while I have been writing up.

Thank you to my darling son Dylan for coming along just to make sure I didn't work too hard! And thank you to baby number two for forcing me to submit on time! Most of all thank you Matt, my husband, for talking things over when I'm stuck for ideas, and taking Dylan to the allotment so I can work, and a million other little things that I can't even put into words.

Publications and Awards

Main Publications

1. Rose Johnson, Nadia Bianchi-Berthouze, Yvonne Rogers, and Janet van der Linden. 2013. Embracing calibration in body sensing: using self-tweaking to enhance ownership and performance. In Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing (UbiComp '13). ACM, 811-820.

2. Rose Johnson, Yvonne Rogers, Janet van der Linden, and Nadia Bianchi-Berthouze. 2012. Being in the thick of in-the-wild studies: the challenges and insights of researcher participation. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12). ACM, 1135-1144.

3. Janet van der Linden, Rose Johnson, Jon Bird, Yvonne Rogers, and Erwin Schoonderwaldt. 2011. Buzzing to play: lessons learned from an in the wild study of real-time vibrotactile feedback. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11). ACM, 533-542.

4. Janet van der Linden, Erwin Schoonderwaldt, Jon Bird and Rose Johnson. 2011. MusicJacket - combining motion capture and vibrotactile feedback to teach violin bowing. IEEE Transactions on Instrumentation and Measurement, 60(1), 104-113.

Papers 3 and 4 report to initial work I carried out in the early stages of my PhD studying the MusicJacket. Due to constraints of space the full reports of these studies are in the appendix. The findings reported in these papers are referenced to as the basis for the real-time feedback designs in this thesis.

Paper 2 reports some of the findings from the studies in chapters 5 and 6 and contains a similar discussion of methodology to that found in chapter 10.

Paper 1 reports some of the findings from the study in chapter 7, and uses these to produce the guidelines for designing calibration discussed in chapter 9.

Awards

EPSRC Pioneers Competition Finalist 2011

The MusicJacket won best demo at Pervasive 2010

Contents

Chapter 1 – Introduction

1.1 Background

Computing has moved beyond the office with the traditional keyboard and mouse interface to cover a wide variety of interaction methods and applications: from smart phones with touch screens and voice recognition, to smart buildings where people's day-to-day movements are enough to initiate interaction (Rogers, Hazlewood, Marshall, Dalton, & Hertrich, 2010). Not only are input methods changing but the way computers provide feedback is also becoming more varied: from vibrations to help users input text on a touch screen (Hoggan, Brewster, & Johnston, 2008), to moving physical installations (Rogers, Hazlewood, Marshall, Dalton, & Hertrich, 2010), to immersive theatre experiences which use both the facilitation and deprivation of sensory modalities to enhance the experience (van der Linden et al. 2011). Visual, aural and tactile senses are now all regularly used for human computer interaction.

One input method that has become increasingly popular in recent years is body movement. Games consoles such as the Nintendo Wii (Nintendo Wii, 2006) and the Xbox Kinect (XBox Kinect, 2010) use motion capture technology such as inertial sensors (Wii), cameras (Kinect) and depth cameras (Kinect) to enable players to interact with the game using their whole body. Motion capture technology is also being researched for non-gaming applications such as physical therapy (e.g. (Geurts et al., 2011)) and motor learning (e.g. (Lieberman & Breazeal, 2007)). Here motion capture technology can be used to measure and categorise movement and give feedback to users about their movements to help them improve the way they use their body. Not only is it possible to give this feedback after the event to help people reflect on their movement, it is also possible to give feedback in real-time while the movements are taking place. This could potentially transform the way we learn motor skills because it would enable people to correct their mistakes before they complete an action. For those that subscribe to the concept of muscle memory (e.g. (Kempter, 2003)), in which skilled movement is learned and proceduralised through repeatedly practising a movement correctly, real-time feedback about movement is the ideal way to enhance training because it can ensure that practice is always carried out correctly.

However this relies on the assumption that people can take in and respond accurately to feedback in real-time. There is a strong argument why this should be the case: we continuously depend on real-time feedback from our senses to negotiate our environment. As such we are well practised in processing complex sensory information in real-time. Moreover, evidence from sensory substitution studies (Bach-y-Rita, Collins, Saunders, White, & Scadden, 1969) suggests that there is plasticity in our sensory experience enabling people to learn to respond to complex real-time feedback that is different from normal forms of sensory perception. Studies which have tested real-time feedback systems to teach movement (e.g. (Bloomfield & Badler, 2008)) have found positive results which suggest that real-time feedback can enhance speed and accuracy of learning specific movements. However, this does not give a full picture of how real-time feedback might be applied to learning a long-term skill such as a musical instrument or a sport. In current research the real-time feedback is the participant's main focus and they are able to consciously respond to it with few distractions. When people learn a skill there are often several elements of the skill that they must perform at once meaning they are not able to give full attention to the real-time feedback. Often the movement is a means to an end rather than an aim in itself, e.g. violinists move to make music and tennis players move to hit the ball; they need to perform the movement correctly but the movement is not their only focus. This may reduce people's ability to use real-time feedback to enhance their learning. Therefore, designing real-time feedback for learning a skill like the violin raises new research questions beyond those related to simply teaching movement. Questions that need addressing are how much feedback people can take in while still successfully carrying out the skill, and what is the best way to deliver the feedback so that it catches attention without being a distraction. This thesis investigates how real-time feedback can be applied to learning to play the violin and through this examines how real-time feedback can be used to support people in learning a complex and long-term motor skill.

Learning to play an instrument like the violin is a difficult task which requires regular and sustained practice for at least 10 years (Ericsson, Krampe, & Tesch-Römer, 1993). Bowing alone has been shown to take 700 practice hours (Konczak, vander Velden, & Jaeger, 2009). Physical movement is considered to be one of the most important aspects of learning to play the violin; indeed some teaching methods would say that it is *the* most important ((Garson, 1973) quoting Suzuki, (Kempter, 2003)). Yet it is perhaps also one of the most difficult parts for a teacher to communicate to their student especially as western teaching increasingly moves away from physical contact with students. A pupil copying the teacher's playing may have problems due to the difference in viewpoint when sensing their own movements compared to watching somebody else's. Even if a teacher does physically move a pupil into the correct stance and movement, the pupil may not be able to repeat this since they have not used their own muscles to make the movement and therefore will not develop the same kind of muscle memory.

Real-time feedback may be able to bridge this gap because it guides players' movements in a different way. By giving near instant feedback as to whether the pupil is moving correctly, real-time feedback can set up a strong link between cause and effect so that pupils are able to understand through their own physical movement and experimentation which movements are correct and which are not. Moreover, there are many sensory modalities through which feedback could be delivered which may enhance this process. One of the most promising is vibrotactile feedback which can be positioned on the body to enhance the link between movement and feedback.

1.2 Research Objectives

The aim of this PhD is to study the way that real-time feedback can be used to aid training in motor skills through studying the practical example of violin. Learning the violin is a highly demanding skill, both in terms of the number of practice hours required to achieve mastery (approx. 10,000 hours (Ericsson, Krampe, & Tesch-Römer, 1993)) and the real-time sensory and cognitive demands of playing. This makes it interesting from the point of view of studying real-time feedback because it challenges us to design feedback which is optimised for use when under cognitive load and enables us to study the limits at which feedback becomes overloading. Learning the violin also involves independent practice, taught lessons and playing together which enables the study of real-time feedback in different settings with different physical and social demands. Similar types of learning situation are found in many other application areas such as sports training which may take place individually or as a group and with or without a coach. Learning the violin is also a long-term goal where factors such as motivation and sustainable learning become very important. This is also relevant to many other possible applications for real-time feedback, for example physical therapy, learning a sport or improving posture at work.

Through studying real-time feedback for learning the violin this thesis aims to contribute to knowledge in three different ways:

1. A theoretical framework for understanding real-time feedback for learning motor skills in-the-wild
2. A set of practical design guidelines for designing real-time feedback for motor learning
3. A methodological contribution in the form of a framework for thinking about the role of the researcher in in-the-wild studies

Research investigating real-time feedback for learning motor skills has shown promising results (e.g. (Bloomfield & Badler, 2008), (Lieberman & Breazeal, 2007)). However, these studies have mainly taken place in the laboratory and have focused on short-term improvements in often quite arbitrary tasks such as mimicking arm posture (Lieberman & Breazeal, 2007). This is very different to what is involved in acquiring a long term skill such as violin playing or any other type of music or sports training. Research which has applied real-time feedback to music learning (e.g. (Ng, Weyde, Larkin, Neubarth, Koerselman, & Ong, 2007)) has tended to focus on the technological challenge, demonstrating what aspects of playing real-time feedback could be applied to. So far, very little research been carried out to ascertain how players would use real-time feedback and how best to deliver information to players as they play. This thesis aims to address this gap in the research by studying real-time feedback for violinists in-the-wild by building prototype systems and studying them with learners in naturalistic settings, looking at both short-term and long-term effects (from hours to months). An in-the-wild study aims to study technology in as natural setting as possible in order to learn about the demands and opportunities that participants will have when using the technology in the real-world. In this case, this means studying real-time feedback with genuine players in realistic settings such as at home or school or other place where they are motivated to play. Our initial work studying real-time feedback for learning the violin, which can be found in Appendix B, tested real-time vibrotactile feedback in the laboratory using an experimentally informed methodology. This demonstrated the potential of real-time feedback as a practice aid, but also revealed pronounced differences between individual players. A second study in Appendix C was carried out in-the-wild with teachers and their pupils. This showed that in the real-world the ways real-time feedback can influence learning are more complex than laboratory studies would suggest. The distractions and cognitive load of the real-world setting and tasks produced different findings to those from the laboratory study. This thesis extends this research by carrying out further in-the-wild studies of real-time feedback. By studying real-time feedback in real-world settings the aim is to move beyond the findings of the laboratory studies which study real-time feedback in isolation and build a new framework for understanding real-time feedback in-the-wild, which is practical and relevant for real-world applications.

The framework is developed based on the findings from a series of user studies (Chapters 5-7) and is discussed in detail in Chapter 8. The process of investigating real-time feedback for motor learning is not only a research process but also a design process. Each study is dependent upon a prototype that meets the requirements of the learners and the study setting, and the findings from each study have practical design implications which feed into the designs in subsequent studies. In other words, an iterative design process takes place throughout the programme of research. The second thread that runs through the thesis culminates in a set of guidelines for the design of real-time feedback presented in Chapter 9.

In-the-wild studies are increasingly prevalent as ways of understanding how new technologies may potentially disrupt, support or enhance our everyday activities. In this paper we use the term *in-the-wild* as Rogers (2011) uses it to mean studies which involve deploying new technologies in real-use, real-world situations and studying how they are used in this context often with the intention of improving a design. Implicit within this type of methodology is the idea that physical and social context will have a critical effect on usage. Part of this social context is the experimental context itself and the role that the researcher plays in the study (Brown, Reeves, & Sherwood, 2011). Traditionally, in laboratory experiments the researcher‘s role is viewed more as a distant, controlling figure in order to test hypotheses using the scientific method. However, the researcher can no longer take such a prescriptive role if the participants are to interact with the technology and environment in a natural way because this will interfere with events unfolding normally in that setting. Moreover, by distancing themselves from the study, the researcher may have difficulty getting an inside view on how the participants perceive the setting and the study. This thesis studies various roles a researcher can take in in-the-wild studies of violin playing, from observer to teacher to participant. This enables reflection on the different roles a researcher might play in-the-wild and how these affect the way a study unfolds. This is discussed in more detail in chapter 10.

1.3 Research Questions

The overarching research question this PhD addresses is:

How can real-time feedback aid learning to play the violin?

This is a broad exploratory question: real-time feedback can take many forms and not all of these can be explored in this thesis. First, feedback can be given using different sensory modalities, such as tactile, kinaesthetic, visual or auditory modalities. In terms of modality this PhD limits itself to investigating the two most promising candidates from the literature review (chapters 2 and 3) namely vibrotactile and ambient visual feedback. These are directly compared in a study in chapter 5 and then employed in complimentary combinations in chapters 6 and 7. A second variable is the amount of information that real-time feedback gives to learners and how this is presented. This will affect how players are able to use feedback while they play, and whether they ignore it or find it useful, distracting or overloading. This is investigated throughout this thesis by studying different designs of feedback in different learning settings. The findings relating to this are then discussed in depth as part of the theoretical framework for understanding real-time feedback presented in Chapter 8.

Learning the violin can also take many forms and therefore further sub-questions arise from the different ways real-time feedback could aid learning depending on the setting and form of learning taking place. In a laboratory setting it is easy to view learning the violin simply as improvements in physical playing technique. However, in the real world learning is a more complex multifaceted process. One form of learning takes place in the lesson, where the teacher communicates with the pupil the way the violin should be played and how they can improve and the pupil tries to replicate what their teacher shows them. In this setting feedback could potentially aid learning in other ways, for example helping teachers and pupils to communicate or reinforcing what the teacher is saying. Another form of learning happens when the pupil practises independently. Here, as well as correcting movement, real-time feedback could aid learning by motivating students to practise or to focus their practice on the goals the teacher has set them. Violinists also play and practise in ensembles. In this situation, the learning focus is not only on individual improvement but also on working together and improving as a group and thus the role real-time feedback may play in aiding learning will be different again. This demonstrates the different roles real-time feedback may play in different settings. Therefore, in order to understand how real-time feedback can be an aid to violin learning, this thesis investigates the effect of real-time feedback in-the-wild in different practice settings. This builds on our previous work which has investigated real-time feedback in a teaching setting (van der Linden, Johnson, Bird, Rogers, & Schoonderwaldt, 2011) (Appendix C). Therefore this thesis aims to investigate three specific research questions:

(i) How do visual and tactile modalities of real-time feedback differ in supporting individual violin practice?

(ii) How can both shared and personal real-time feedback support playing as part of an ensemble?

(iii) How can real-time feedback fit into the extended process of learning through various forms of practice?

Question (i) is addressed quantitatively and qualitatively by a study of individual practice in Chapter 5. Question (ii) is addressed qualitatively in an in-the-wild user study held at a week-long musical summer school in Chapter 6. Question (iii) is addressed in Chapter 7 in a 14 week in-the-wild study held at a high school which investigated both ensemble and individual home practice. These three studies also investigate different user groups, the first two are with adult amateur players, the last study is with high school children. Through studying real-time feedback in these multiple settings the aim is to understand the role real-time feedback can play in aiding learning to play the violin in a naturalistic and comprehensive way which takes into account the importance of human factors such as social relationships and motivation in the way people learn.

Learning can also happen over different periods of time; real-time feedback may act as a quick fix for a minor problem or may need to be used in the long term to train a player into a new way of moving their body when playing. New technology also has a novelty value which may make players particularly motivated to use it early on but this enthusiasm may wane over time. On the other hand, as players become more familiar with the feedback they may be able to use it more effectively and appropriate it to their own approach to playing. Therefore, it is important to study the effects of real-time feedback in both the short and long term. The study in Chapter 5 focuses on the effect of real-time feedback in a single practice session. The study in Chapter 6 extends this to studying an intensive week of practice at a summer school. Finally Chapter 7 studies real-time feedback over a much longer period of 14 weeks.

1.4 Thesis Structure

This thesis comprises 11 chapters summarised in Table 1-1. Chapters 2-4 cover literature reviews, methodology and preliminary work. Chapters 5-7 present three empirical studies investigating real-time feedback. Chapters 8-11 discuss the findings from these studies and draw conclusions.

Chapter 2 reviews the current research which uses real-time feedback to enable motor learning and music learning and in other domains such as encouraging good health and fitness, physical therapy, navigation and sensory substitution. This review aims to draw lessons from the research to inform the design of real-time feedback for learning the violin. Chapter 3 then takes a more theoretical approach by studying relevant theory and experimental evidence from psychology, sports science and human computer interaction. The main areas for review are attention, modality, theories that link perception and action and theories of motivation. This literature is used to draw out implications for the design of real-time feedback for the violin and to form a basis on which to later build a theoretical framework to describe real-time feedback (Chapter 8). This chapter also summarises our previous research about real-time feedback for learning the violin ((van der Linden, Schoonderwaldt, Bird, & Johnson, 2011), (van der Linden, Johnson, Bird, Rogers, & Schoonderwaldt, 2011)). Chapter 4 presents the research methodology used throughout the thesis. Chapter 5 compares ambient visual feedback or vibrotactile feedback in a user study, as these appear from the literature to be the most suitable modalities for aiding violin playing. This study was conducted with orchestra players and this raised further research questions about whether shared real-time feedback could be used in ensemble playing. This is investigated in Chapter 6 in an in-the-wild user study held at a musical summer school. Chapter 7 presents the final user study in this thesis. This was held with participants from a high school orchestra and aimed to investigate real-time feedback in all forms of practice: ensemble playing, taught playing and home practice.

Chapter 8 takes the findings from these three individual studies and forms them into a theoretical framework for understanding real-time feedback in-the-wild. Chapter 9 uses this framework and the findings to formulate practical design guidelines to inform the design of real-time feedback for motor-learning. Chapter 10 reflects on the methodological lessons learned over the course of the research in this thesis. In particular, it focuses on the researcher's role in in-the-wild studies and formulates a framework for analysing this. Finally, Chapter 11 draws these three discussions together and summarises the key conclusions from this investigation into real-time feedback.

Ch	Title	Summary
2	Feedback Technology	Literature review of current technology that gives real-time feedback to aid learning music, sports, motor skills and to replace or augment the senses.
3	Attention, Modality Perception and Action and Motivation	Literature review of experimental findings and theoretical concepts relating to how people direct attention, the properties of different sensory modalities, the relationship between perception and action and motivation. This chapter finishes with a description of the early studies carried out in my research of real-time feedback using a prototype call the MusicJacket.
4	Methodology	This chapter describes the in-the-wild methodology used in this thesis and how the analysis has moved from a quantitative to a qualitative one.
5	MuSense: Comparing Visual and Vibrotactile Feedback	This describes a user study of real-time feedback for improving bow length in individual practice which compares three ways of giving real-time feedback: vibrotactile feedback, and two different versions of visual feedback.
6	ShareSense: Studying shared and individual feedback in ensemble playing	This chapter describes a study of real-time feedback for improving bow length at a musical summer school investigating how shared feedback can be used in ensemble playing.
7	Twinkly Lights and Buzzy Jacket: Multimodal feedback as part of a programme of learning	This chapter describes a 14 week study of real-time feedback for improving posture in violin players at a high school, covering ensemble playing and individual practice.
8	Framework to Describe Real-time Feedback for Motor Learning In-the-Wild	This chapter describes a framework of how real-time feedback is used and the factors which influences this.
9	Guidelines to Design Peripheral Real-time Visual or Vibrotactile Feedback	This chapter builds on the framework to produce a set of practical guidelines for the design of real-time feedback.
10	Reflections on Methodology	This chapter reflects on some of the challenges and opportunities that come from using an in-the-wild methodology and how the outcome of a study is influenced by the role the researcher chooses to take.
11	Conclusion	This chapter outline the key conclusions of the research in this thesis.

Table 1-1: Thesis Structure

Chapter 2 – Feedback Technology

2.1 Introduction

Technology has been used to give real-time feedback in many different application areas from music learning to navigation. In particular, studies of real-time feedback across a number of application areas have shown that it can play different roles depending on the application: guiding movement, alerting people to mistakes, motivating people, or replacing or augmenting senses. Although this PhD is focused on how real-time feedback can be applied to music training, previous research which investigates how feedback can enhance other areas of human behaviour and performance can provide a deeper understanding from which to understand the specific domain of learning to play the violin. Therefore, this literature review will cover how technology can be designed and used for helping people learn a range of skills, including music, motor skills, improving posture and fitness, and rehabilitation. It also draws on research about real-time feedback from other applications such as sensory substitution, mobile applications and gaming.

The literature review begins with an overview of studies that have used real-time feedback for helping users when learning to play musical instruments. It then considers real-time feedback for learning motor-skills to gain more insight about how users might use a real-time feedback system to learn and train. In contrast, the chapter then moves on to examine studies about how real-time feedback has been used for posture and fitness in the real-world where participants have physical, cognitive and social demands that affect how they use feedback. The role of real-time feedback for rehabilitation and physical therapy is also examined to understand more fully how technology interventions can motivate people. The last section examines other kinds of real-time feedback applications that are not so focused on movement and motor learning, including real-time feedback to replace or augment senses and real-time feedback to improve text entry on mobile phones.

2.2 Feedback for Learning Musical Instruments

There are several research projects which have looked at giving real-time feedback to musicians to help them to improve their playing. These show the potential of new sensor technologies to measure key aspects of playing technique and give feedback in real time. However, there is often very little user evaluation in these projects making it difficult to learn from them how best to design feedback for musicians.

The iMaestro project (Ng, Weyde, Larkin, Neubarth, Koerselman, & Ong, 2007) investigated a method of giving feedback to musicians by creating a three-dimensional augmented mirror that enabled musicians, mainly string players, to view themselves playing from any angle and record and playback sections. It also displayed other measurements of their playing such as bow speed using dials on screen. The initial user tests indicated that musicians liked the visualisation and were interested to analyse their playing in this way. They found that the musicians noticed different patterns in the way they played which they had not noticed before. However, there are potential drawbacks to such detailed visual feedback. For example, they will not be able to read music whilst using the feedback in real-time. As an alternative method of feedback the project also looked at sonification of gestures, such as parallel bowing. This means informing the player about their bowing by making a sound to indicate whether they are on a straight trajectory. The sounds either took the form of a bell or a processed version of the sound of the instrument. In informal testing with teachers, some teachers commented that they felt that the sonification might distract learners from the sound of the instrument, whereas others said it may be useful in certain situations, however this has not been evaluated.

Grosshauser and Hermann (2009) have proposed a system which uses auditory feedback built into a music stand as a form of sonification of the player's movements. Rather than the alarm bells used in the iMaestro project, Grosshauser and Hermann propose encoding much more information in the sound of the feedback. For example, adding an addition note to indicate that the player is moving their bow hand in a way that might hit the next string, or changing the balance between two stereo speakers to indicate in which direction the player is deviating from the ideal bowing trajectory. However, there are not any published evaluations of the feedback system so it is unclear as to how effective this form of auditory feedback is. Grosshauser et al. (2012) have applied a similar type of sonification to ballet dancing training. This system received positive feedback both from teachers and students, the teacher observed a difference in the students' movement and the students felt that they could understand and respond to feedback in real time. However, dance training is quite different to music training, and it may not be possible for musicians to attend to, and process, the complex information encoded in the sonification sounds while at the same time listening to the sound of their instrument. There may be a danger that manipulating the sound of the instrument will prevent learners from developing an understanding how their movement affects the sound of the instrument itself and could prevent them from developing their own individual style.

Another possible modality for giving real-time feedback that has been investigated is vibrotactile feedback. Grosshauser and Hermann (2009) attached a variety of sensors to the violin and bow which allow the precise measurement of bowing movements and also give vibrotactile feedback to the student on aspects of their playing. They mention that teaching with this type of feedback can allow the teacher to provide 'silent hints' that are physical, rather than verbal. They propose that this will be less disturbing for the pupil as they play. However, there is no systematic evaluation of how such hints are experienced by the pupil or for which type of activity they are most useful.

Besides real-time feedback there have been other technology interventions developed which aim to help people learn musical instruments. One example is MirrorFugue (Xiao & Ishii, 2011) which was designed to help people learn to play the piano from a teacher remotely. It projects the hands of the teacher on a vertical display surface just in front of the pupil’s piano keys. This was compared to two other conditions: dots positioned above the keys to indicate which ones to play, and a remote video link with no addition technology. Xiao and Ishii (Xiao & Ishii, 2011) found that the projected hands were the most popular interface with learners. They were also jointly the most effective, alongside the dots interface, in terms of the speed at which participants learned to play a nursery rhyme. This shows that positioning information in a way that is meaningful in terms of how the body is used to play the instrument can enhance learning.

Holland et al. (2010) built a system which uses rhythmic vibrations on the body to teach drumming. Rather than giving feedback it aims to help people learn to play polyrhythms on a drum kit by feeling the beat on the relevant parts of their body that they use to play the different drums. A small scale user study was conducted with five participants who all experienced learning complex rhythms in three different conditions: audio only, vibrotactile only and audio and vibrotactile combined. The audio signals took the form of drumbeat guide tracks heard through a stereo, the vibrotactile signals mirrored the audio tracks but were spread across the body on the limbs which corresponded to the correct drum. All the participants preferred the combination of audio and vibrotactile signals commenting that the vibrotactile signals helped them to feel the beat with their body by showing which limb was needed to play which beat, which the audio did not do. However, for fast rhythms it was difficult to feel exactly when the beat started with the vibrations. Holland et al. also observed that in the vibrotactile conditions, participants learnt the rhythms by playing with multiple limbs simultaneously, whereas in the audio only condition participants learnt by practising rhythms with each individual limb at a time. This study shows some of the potential advantages of using vibrations to communicate body movement. Like the projected hands in the MirrorFugue interface, the audio and vibration combination presents information in a way that makes a link between music and the way the body is used to create it. The way vibrations are positioned and felt on the body shows what movements are needed in a clear and direct way and enabled participants to engage with the rhythm using their body, whereas the audio feedback was not so embodied.

Another study which used vibrations to help people to learn to play music was conducted by Huang et al. (2010) and investigated passive vibrotactile learning. Participants learnt a simple tune on the piano which only used five notes, mapping each finger directly to a note on the piano. They then carried out a distraction task of reading comprehension for 30 minutes. During this time half the participants wore a glove which vibrated the sequence of notes they had learnt while also listening to the tune. The other half of the participants were played the tune without the vibrotactile stimuli. The participants were then tested on their ability to play the tune. The group who had the vibrotactile stimuli performed significantly better than those who only had the audio. Moreover, on average the participants with the vibrotactile stimuli improved, whereas the performance of those without the vibrations deteriorated. This finding shows that the tactile modality can be used to facilitate passive learning of sequences of movement when people are not actively attending to it. This suggests that vibrotactile feedback might be well suited to giving real-time feedback to violin players because in this case the feedback needs to be attended to only peripherally. However, these results are limited in scope because in this study the mapping was one finger to one note whereas when playing a musical instrument beyond the very basic level the movements are much more complicated and may not easily be encoded in vibrotactile sequences that that can be learned passively.

Huang et al. (2010) also conducted a follow-up study comparing passive and active learning with unskilled participants and skilled musicians. In the active condition participants tried to replicate a sequence of notes that were lit up on the keyboard, and repeatedly tried again until it was correct. In the passive condition, each time the participants failed an attempt at playing the sequence they went through a passive training phase where they carried out a distraction task and used the vibrotactile glove. The unskilled learners took significantly fewer attempts at learning the sequence in the passive training condition compared to the active condition. Conversely, the skilled group took fewer attempts in the active learning condition and in the passive condition many of the skilled musicians had to abandon learning the sequence because it was taking so long to learn. Huang et al. (2010) suggest that this difference comes from the sequences of notes being amusical and that while the unskilled learners could passively learn any sequence, for the skilled musicians it needed active concentration to be able to overcome their training which biases them towards playing sequences that make musical sense. These findings suggest that once a habit has been formed passive learning may not be a strong enough method to overcome it. These results also make it questionable whether passive learning has a role to play in music training because learning a musical instrument is much more than learning a sequence of movements, it is about learning how music and movement are interconnected so that the musician can apply this knowledge creatively and expressively. It may be that active engagement is necessary to achieve this kind of knowledge base. I would certainly argue that the pleasurable part of music making comes through active engagement in it.

Another form of real-time feedback for violin learning is force feedback. Baillie et al. (2005) used a PHANToM haptic feedback arm to simulate the physical sensations involved in bowing the violin. This is a robotic arm positioned on a table which players hold and move back and forth like a bow on a violin string. The PHANToM arm is able to simulate haptic sensations of stiffness, friction and damping that come from the bow slipping against a violin string. To aid learning artificial constraints were also added, such as force feedback which makes it difficult to move the 'bow' up and down the string to help people to learn to play with the bow positioned correctly on the string. This prototype was tested with seven violin players and results suggested that the prototype was effective in simulating the sensation of bowing. The suggestion of artificial constraints to guide bowing was also positively received.

However, there are drawbacks to this technology for giving real-time feedback. Using a PHANToM robotic arm to give force feedback does not easily fit into normal practice with a real instrument. However well the PHANToM may be able to simulate bowing, it will be a different experience to playing and holding a real violin and learning about how the physical action and sensation involved in bowing relates to the change in the sound that the violin makes. Therefore, force feedback implemented in this way could only play a limited role in learning the violin. Another method of giving force feedback would be to use a wearable exoskeleton which attaches to the player's arm to push their limbs. For example, a robotic arm has been attached to the arm of stroke patients to guide them in performing reaching tasks more accurately during rehabilitation (Frisoli, Salsedo, Bergamasco, Rossi, & Carboncini, 2009). This will allow players to have force feedback while playing a standard violin, however this technology is currently bulky to implement which may hinder freedom of movement when playing. Other methods for force feedback also include jets of air which can push against the body. For example, Suzuki and Kobayashi (2005) used an array of air jets to give force feedback that corresponds to interacting with objects in virtual reality. However, participants could not interact directly with the objects, instead they had to hold a spoon shaped air receiver, and feel the objects through this. This was so that they did not feel the sensation of the air again their skin.

Besides difficulties with implementation, there is a more fundamental reason against choosing force feedback as a way of aiding learning motor skills. Research has shown that passive movement where the body is moved for someone does not develop the same kind of muscle memory as active movement where a person voluntarily moves their limbs. For example, Lotze et al. (2003) compared participants moving their wrist voluntarily in response to a visual signal, with participants having their wrist moved passively by a motor. The aim was for participants to learn to do wrist movements of a particular duration. The participants performed significantly better after active training with voluntary movements than those who had the passive training. fMRI results also show different kinds of cortical reorganisation in active movement compared to passive movement. In particular, Lotze et al. found significantly greater levels of activation in the M1 region of the brain with active movement than with passive movement. This area of the brain includes the primary motor cortex and is important in motor learning. This suggests that giving feedback which encourages players to actively correct their movements should lead to faster learning than force feedback which moves a player's limbs into the correct position. For this reason, this thesis focuses on feedback such as vibrotactile feedback which requires participants to make voluntary movements to correct their playing.

The technology interventions in the field of music learning show that systems can be built that sense the key movements of players' arms and hands that can provide guidance on when best to give real-time feedback. The studies reported show how choice of modality is important as well as presenting information in a way that references how the body is used to play the instrument. Of all the feedback modalities, vibrotactile feedback appears to have the most potential for guiding movement. However, it is hard to determine the extent of efficacy as there are no reported user studies that have investigated the benefits of using real-time vibrotactile feedback for helping musicians play. In the next section literature is reviewed from other domains where real-time vibrotactile and visual feedback to enhance motor learning has been studied in more depth.

2.3 Feedback for Learning Motor Skills in Other Domains

Studies of real-time feedback to enhance motor learning have shown that real-time vibrotactile can be particularly effective for learning movements compared to visual feedback alone. Lieberman and Breazeal (2007) built TIKL, a system of vibration motors on the arm which aim to guide movement. This used both single vibrators and the technique of sensory saltation (Geldard & Sherrick, 1972) patterns to communicate movements to the participants wearing it. Sensory saltation is an illusion where patterns of vibrotactile bursts spread over two or more vibrators can give the impression of vibrations running up or down a limb. The effect was originally discovered using taps on the skin with solenoids rather than vibrations and was referred to by its discoverer as the "cutaneous rabbit" because it felt like a rabbit was hopping up his arm. TIKL was tested in a between participants user study with 40 participants where participants had to mimic the arm position and movements of a person on screen. TIKL was used to give feedback about how close the participants were to the desired position. Half the participants used TIKL and had visual feedback on screen; the other half of the participants only had the visual feedback. Participant using TIKL performed significantly better in mimicking the arm position and movements. The vibrotactile feedback was particularly good at guiding the angle of the joints (e.g. the angle of bend of the elbow or wrist) whereas it had a smaller effect on correcting the rotation of the joints (e.g. twisting the arm or hand). Lieberman and Breazeal suggest that this is because the feedback about joint angle is easier to understand because it is a simple vibration on the part of the body that is out of alignment that responds directly to arm being in the wrong position – they liken this to a "force-field" around the correct movement; whereas the rotation information uses the saltation patterns to indicate the motion of the rotation that needs to be made. This suggests that simple feedback using single vibrators rather than patterns of vibrators may be easier to understand; however, the difficulty with the rotation could also be due to rotation movement itself being a more difficult type movement to understand rather than a problem with the feedback.

Bloomfield and Badler (2008) also built a vibrotactile array to guide movement by giving real-time feedback. They conducted two experiments with it. In the first one, participants had to reach into virtual shapes while avoiding colliding with them. Collisions were indicated either with a visual or tactile alert or a combination of both. The groups with the tactile feedback collided significantly less than those without the tactile feedback. Surprisingly, those with only the tactile feedback performed better than those with the combined feedback. Bloomfield and Badler suggest that this is because the combined visual and vibrotactile alerts gave too much information, with the visual alerts distracting the participants from the more effective vibrotactile alerts. In the second experiment, five participants were taught a series of karate moves. These were during a series of training sessions, firstly by following visual instructions followed by being provided with tactile feedback. Participants improved significantly during the first vibrotactile training session, suggesting that the vibrotactile feedback was helping them to learn how to practise the movements by giving embodied feedback directed to their movement which helps them to feel which limb to move and where they were making mistakes. However, a limitation of the study is that all the participants used the vibrotactile feedback so there is no comparison condition to tell whether the improvement was due to having feedback about movement, or specifically having real-time vibrotactile feedback, or whether this is a learning effect from having some time to train irrespective of the feedback.

These studies show that vibrotactile feedback can be effective for guiding movements. This raises further questions about where vibrators should be activated on the body to provide signals to people that are easy to understand in terms of where they are placed on the body and how they appear to move. Spelmezan et al. (2009) investigated this by positioning various vibrator arrays on the thighs, torso shoulders and chest. They conducted a study where participants were given different vibrotactile patterns on the body and asked to respond with whichever movement they felt was most appropriate. Their findings show some patterns in people's responses but a lot of individual differences as well. They found that when they gave pulses of vibration on a particular part of the body participants generally responded by moving that part of the body in some way, but there were individual differences in the type of movement. They found that directional patterns such as vibrations running down the back of the thighs gave slightly more consistent responses; in the case of the pattern on the back of the thighs 60% of people bent their legs. More complex patterns were found to produce similar responses to directional patterns. Spelmezan et al. observed that half the participants tended to move away from the vibrations and the other half tended to move towards the vibrations.

This study shows that although the placement and pattern of vibrators can suggest a particular type of response, the exact response is not consistent across all people and therefore vibrotactile feedback must be learned rather than the same patterns being instinctively felt by all. Spelmezan et al. (2009) investigated how easy vibrotactile instructions are to learn. They devised their instruction set using a push metaphor, where participants were asked to move away from the vibrations. They also used directional vibration patterns to indicate the direction of movement. This draws on the findings of the previous study by choosing some of the most popular responses to the feedback from the initial study but arranging them so that the feedback works to a consistent metaphor to make it easier for people to learn. Participants were given ten minutes to learn a set of ten vibrotactile instructions. They were then tested on them, first in a relaxed setting, then while playing a snowboarding game on a Nintendo Wii balance board. In these tests they had to respond to randomly triggered vibrotactile instructions by saying the correct response and carrying out the movement. In both conditions participants performed extremely well, with some movements being performed correctly by 100% of participants even under the cognitive and physical load of playing the game. The average correctly performed movements were above 90% even on the balance board condition. This shows that tactile instructions can be learnt quickly in situations where participants are able to give attention to the tactile feedback.

Spelmezan et al. (2009) used the same system of instructions and tested on people going down a real ski-slope. They showed that participants could learn to respond to these instructions with an accuracy of 87%, even when snowboarding down a slope. This is less accurate than those who were given audio instruction. Furthermore, the reaction times were much faster for the tactile feedback condition as the participants could react as soon as they felt the vibrations rather than having to wait for the whole verbal command. In addition, the findings from a questionnaire showed that participants found the tactile instructions were less distracting although they thought verbal instructions were easier to map onto the movement.

McDaniel et al. (2012) have also investigated how to position vibrators in an intuitive way. Instead of using a push metaphor, they chose to use the concept of *"follow-me"* to indicate the direction of motion: vibrations moving up the arm suggest that the arm is being pulled up; when they move down the arm, the arm is being pushed down. Similar direction patterns are used around the arm and wrist to indicate rotation. In addition, they aimed to communicate information about the speed of movement through pulsing the vibrators at different speeds. Faster pulsing indicates the need to move more quickly, while slower pulses indicated that participants should slow down the movement. They conducted a study to investigate the *'distinctness, naturalness and usability'* of this set-up. Sixteen participants went through a familiarisation and training stage where they learned the movements that corresponded to each vibrotactile signal. This training stage was complete once each participant had reached an accuracy above 80%, which took on average running through the feedback instructions 1.25 times for positioning and 1.12 times for speed. Next was the test stage where they experienced the vibrotactile feedback instructions in a random order and had to respond with the correct movement. Participants recognised the position feedback correctly 94% of the time and responded immediately 91% of the time; participants recognised the speed feedback correctly 90%. Participants also rated the ease of recognition, learning and intuitiveness of the vibrotactile instructions highly, with most of the movements scoring above 4 on a 5 point Likert scale, where high scores indicated a positive attitude towards the instructions. However, it should be noted that these results are from a study where participants focused solely on the vibrotactile feedback. In situations where participants have to focus on something else while using the feedback, which will be the case in many motor learning situations, this level of complex feedback may become harder to respond to.

Spelmezan (2012) has since conducted a further study investigating how vibrotactile motion instructions can be used practically by learner snowboarders on the slopes and found that the demands of the task affects how participants are able to respond to tactile instructions. He conducted a study where ten participants learnt snowboarding techniques under two different conditions: a normal lesson with an instructor and a normal lesson with an instructor plus tactile instructions delivered while the participant went down the slope each time they changed direction. The vibrotactile instructions consisted of vibrations running down the thigh, intended to remind participants to shift their weight onto their forward leg in the direction that the snowboard was going downhill, and vibrations circling the shoulders to remind participants to turn their body correctly as they turned the snowboard. Participants experienced both conditions, and the order they experienced the feedback was counterbalanced for order effects. The study showed no main effect for vibrotactile instruction on performance, but an interesting interaction effect was found: the tactile instructions were more effective in the afternoon sessions. Spelmezan suggests that this is because participants were more familiar with the techniques by the afternoon sessions and so were practising something they knew, rather than learning something new. This is supported by the comments of the participants who said they struggled to pay attention to the tactile instructions when learning a new technique because they were focusing on the verbal instructions they had been given by their teacher; whereas those practising old techniques did not find the vibrotactile instructions difficult to attend to. This shows how important it is to take into account the demands of the task when investigating tactile feedback. Spelmezan also found that the tactile feedback was more effective for improving upper body posture than changing the weight distribution between the legs. He hypothesises that this might be because distributing weight onto the forward leg is a more difficult technique than correcting upper body posture. He also suggests that giving feedback about two elements of technique simultaneously may be too demanding. It should also be considered that perhaps moving weight onto the forward leg is not only a leg movement but a whole body movement and that the vibrotactile feedback may have been more effective if it had been distributed across the body as well as the leg.

In sum, this review of motor learning research shows that vibrotactile feedback can be effective at providing tactile instructions on the body in real time in a way that can enhance learning compared to visual or verbal guidance alone. However, it also shows that there is no single intuitive layout for the position and pattern of vibrotactile feedback. A variety of signals can work as long as the metaphor used underlying the tactile pattern is kept consistent and participants are given time to learn it. Many of the systems described here give quite complex sets of vibrotactile instructions, and the evidence is that people can respond to these with a high level of accuracy in a lab setting. However, when vibrotactile instructions are applied to learning real physical activities the evidence from Spelmezan (2012) suggests that vibrotactile instructions should be limited to a much smaller set of movements as the demands of the task can become overwhelming. Finally, vibrotactile instruction seems to be most suited to practising skills that learners have already started learning and have some understanding of, rather than giving vibrotactile instructions while learners are trying to learn a new technique from their teacher.

2.4 Feedback for Posture and Fitness

Another area where technology has been developed to provide real-time feedback is health and well-being. In particular there are many applications designed to encourage office workers to sit with better posture or carry out exercises while at work to avoid back problems. The systems do this by alerting workers when they are sitting with bad posture or have remained sedentary for an extended period. Both visual and vibrotactile feedback have been provided to give these alerts. Within the visual modality some of the feedback maybe classed as being in the users focal visual field, for example on their computer screen. Other types of visual feedback are more ambient, for example a plant appearing to sag on a person’s desk indicating they too are slouching.

One of the earliest studies investigating real-time feedback to improve posture was carried out by O’Brien and Azrin (1970). In this study they attached a sensor to the back of the participants in the form of a mechanical switch that was pulled open by rounded shoulders to sense when they were slouching. This was attached to a vibrator which gave participants vibrotactile feedback whenever they slouched. O'Brien and Azrin conducted a within subjects lab study in which eight participants wore the apparatus for four hours, with the feedback being active for the middle two hours and inactive for the first and last hours. A significant difference was found in the amount of time spent slouching between feedback conditions and the non-feedback conditions. Time spent slouching reduced from 54% to 8% between the first two hours (no feedback → feedback). This shows that real-time vibrotactile feedback can be effective in changing posture in a laboratory setting. However it does not specify what participants were doing while using the feedback. This study also does not take into account whether such a device could be used in everyday life where there are distractions and other factors such as social context that might override the efficacy of the feedback to alert them to change their posture.

A commercial device called iPosture (Schnapp & Schnapp, 2007) provides vibrotactile feedback indicating when someone is slouching in a similar way to O'Brien and Azrin's system but using an accelerometer to measure the angle of the chest, rather than a switch on the back. We conducted a small scale study where four users wore the iPosture for several days, including in the workplace, and kept a diary of their experiences (Johnson, van der Linden, & Rogers, 2010). It was found that the vibrotactile feedback was effective for making people aware of their posture, even though the sensing mechanism was not always accurate. However, it did not integrate well in the social activities of everyday life; wearers found that the sound made by the vibration and their involuntary response to feeling the vibrations caused them considerable social discomfort. For example, the sound of the vibration in work meetings made one participant feel self-conscious and for another the sensation of the vibrations startled her during the family dinner. Another issue was that although for some people the vibrations were very noticeable in one case to the point of discomfort, for others they could not feel the vibrations and were unsure whether the device was switched on. This demonstrates that there are different challenges when designing real-time feedback for a real-world context compared to testing in a laboratory and that individual differences in the perception of vibration are emphasised by the real-world setting.

Other devices have also been designed to give feedback in the workplace about posture and health. Haller et al. (2011) compared the effectiveness and appropriateness of three types of feedback cue to encourage people working at a desk to carry-out posture exercises if they were sitting with bad posture, or sedentary for too long. The three types of feedback were a graphical alert which came up on the user's computer screen, a physical plant on their desktop which shook as an alert, and a vibrotactile alert built into the chair. A within participants user study was carried out to compare these different types of feedback. Participants were given three different types of computer based task and experienced each type of feedback in turn for each task. Vibrotactile feedback performed best in terms of people responding quickly to alerts. However, the vibrotactile feedback was perceived by participants as interrupting their workflow and being more disruptive. A third of participants said that they might switch it off if they were using it in a long term study. The graphical alerts were not as effective as the vibrotactile feedback but were still perceived as interrupting workflow and being disturbing. On the other hand, the physical plant on the desktop was seen by participants as interrupting workflow less and was not very disturbing; however, people responded less quickly to it. In this situation the plant on the desktop was found to be the most appropriate way of alerting people about their posture because it alerts people to the need for exercise without forcing them to respond immediately. Carrying out the posture exercises is not very time critical so the delays of approximately 30 seconds which occurred when using the plant device will not impact on the effectiveness of the intervention in terms of health. In other applications where the user's response is more time critical the immediacy of the vibrotactile feedback may be more useful.

Zheng and Morrell (2013) have also compared visual and vibrotactile feedback for encouraging good posture in office workers. They built a chair which can sense the sitting posture of the person in it. Built into the chair was vibrotactile feedback under the thighs, lumbar, centre back and shoulders. These were used to indicate to the user how to correct their posture when they were not sitting correctly. They indicated the same information visually using a sidebar on the computer screen with areas of the body highlighted and an icon of a person sitting in different poses depending on the user's posture. The two modalities of feedback were compared using a between participants experimental study. Participants were asked to respond to the feedback while carrying out their normal tasks on the computer for 30 minutes. Both modalities were found to be equally successful in encouraging good posture compared to baseline measurements collected for 15 minutes before the feedback was activated. When the feedback was switched off for five minutes during the middle of the study without participant's knowledge, the participants maintained good posture, showing that this kind of intermittent feedback can be effective.

Zheng and Morrell also conducted a dual-task study where participants used the feedback while speed typing. The speed of typing with and without the posture feedback was measured. Using either modality typing speed decreased when the feedback was switched on. The visual feedback slowed down typing less than the vibrotactile feedback. People also responded slightly faster to the visual feedback compared to the vibrotactile feedback. This is interesting because the multiple resource theory (Wickens, 1980, 2008) and Mayer's theory of multimedia learning (Mayer & Moreno, 1998) both argue that the more sensory channels that information is spread across the better a person is able to process it. This result does not support this theory. However, two factors should be taken into account: first, Wickens suggest that there are two visual channels, focal and ambient and therefore the posture feedback may be ambient and the typing maybe focal; second, although there is some discussion of vibrator positioning, this may not be optimised, whereas the visual feedback used an icon of a stick man which may be easier to understand. Zheng and Morrell also point out that this study was only short term, as people use the feedback more their responses may become more proceduralised and this may affect which modality performs best. Even over the 15 minutes where they were performing the dual task participants' typing speed improved showing that learning does take place.

Another example of ambient visual feedback to encourage fitness and healthy activity is MoveLamp (Fortmann, Stratmann, Boll, Poppinga, & Heuten, 2013). This is a lamp positioned on a worker's desk which changes colour depending on the number of steps taken over the last two hours. Green indicates positive feedback whereas red is negative. Participants walked significantly more when using MoveLamp, taking on average 57% more steps. The participants said that MoveLamp was negligibly distracting, showing the benefit of using an ambient visual display. This example shows the power of positive and negative feedback for motivating people to change their behaviour.

Another application that used visual feedback to motivate people was the UbiFit Garden that uses a glanceable display on a user's mobile phone to encourage physical activity. The metaphor used to give feedback is of a garden growing healthily or decaying, depending on the level of exercise performed by the user. The motivation is to exercise by watching and nurturing your own virtual garden change over time. The application automatically logs activity and can be used as a journal to record and plan fitness activities. Consolvo et al. (2008) conducted a three-month field experiment in which different groups of people used the UbiFit application with or without the glanceable display. They found that having the glanceable display motivated users to exercise, particularly over the holiday period; this resulted in a statistically significant difference in terms of duration between the participants with the display and those without the display.

Both UbiFit Garden and MoveLamp both show an important function for feedback, namely motivating people and rewarding them for achieving goals. It is interesting that both designs have chosen ambient visual feedback rather than another modality. It raises the question whether vibrotactile feedback could also play a motivating role or whether this is something that only visual feedback is suitable for.

2.5 Feedback for Motivation in Rehabilitation and Physical Therapy

Motivation also plays an important role for feedback in physical therapy and rehabilitation applications. Physical therapy exercises are sometimes viewed as boring (Geurts et al., 2011), or if people have experienced pain from exercise in the past they may be frightened of physical activity (Singh et al., 2014). Moreover, in the case of some patients who are recovering from a stroke or a head injury it may be difficult to explain to them why they must do particular activities due to brain damage and so real-time rewards maybe the only way to motivate them. Music and games have both been used to give motivating feedback to patients in physical therapy applications such as rewarding good posture with music (Kearney & Fussey, 1991) and turning physical therapy exercises into video games (Geurts et al., 2011).

Kearney and Fussey (1991) looked at a brain-damaged patient who had difficulty maintaining an upright head posture. They used a headband with a tilt switch built in to measure the angle of the head and then connected this to a personal stereo playing music that he liked. When the patient held his head in a good posture the music would play but when he allowed it to drop the music switched off. As a result of the feedback the time spent in correct head posture increased significantly from 37% to 76%. This shows how giving a person a reward that they find intrinsically pleasurable to reinforce good movement can encourage improvement.

Singh et al. (2014) used a simple musical phrase to encourage people with chronic pain to exercise. In this case it was an ascending scale of five notes to a comfortable stretching position and then a descending scale back down as they stretched further. People found this simple feedback motivating because it showed them how far they were stretching. This is a different kind of reward to the music feedback used by Kearney and Fussy (1991), because the feedback is rewarding because it indicates achievement, not because it is intrinsically pleasurable in itself. This model for giving motivating feedback is also very relevant to music learning because players are also motivated by a sense of achievement, (Smith, 2005) therefore feedback which lets them measure their progress should be motivating. This way of making feedback motivating should not be dependent on the modality of the feedback or the feedback being intrinsically pleasurable, although this may enhance the effect.

Another way of making feedback motivating is through gamification, i.e. playing games. One example is Geurts et al.'s (2011) games for physical therapy which incorporated movements from physical therapy exercises into video games to make them more motivating. They found that calibrating and adapting games to the needs of the individual players was important to give the right level of challenge to make them engaging and to make sure that they encouraged the correct movements. This was important for their user group because they had varying levels of motor disabilities.

2.6 Other Applications of Real-time Feedback

So far, this review has focused on real-time feedback for guiding or encouraging movement. This section reviews findings from a selection of other applications that are relevant to the design of real-time feedback. Studies of sensory substitution (Bach-y-Rita, Collins, Saunders, White, & Scadden, 1969), which involves replacing a missing sense by delivering the same information in another sensory modality (e.g. converting a visual image into a tactile display to replace lost sight) reveal what happens when people use real-time feedback in the long term whereas nearly all the studies described so far focus on short-term improvement. Next, we examine findings from sensory augmentation and navigation studies that reveal interesting individual differences in the way people respond to vibration ((Nagel, Carl, Kringe, Märtin, & König, 2005), (Pielot & Boll, 2010)) and new sensory experiences that come about through real-time feedback ((Neely & Burström, 2006), (van der Linden et al., 2011)). Finally, we look at studies about vibrotactile feedback for text entry on touch screens that show that vibrotactile feedback can alert people to mistakes in their typing through rough and smooth vibrations whereas the majority of literature here uses the position of the vibrator as a communication the method.

2.6.1 Feedback for Sensory Substitution

The motor-learning studies all investigate vibrotactile feedback in the short term: at most a retention test is conducted a day after the instructions have been learnt. Learning the violin is a long-term goal, and the real-time feedback will be used over a much longer period than a single day. In order to understand more about the effects of vibrotactile feedback on learning over the long term we look at the topic of sensory substitution. By this is meant substituting one sense with another, for example teaching people to perceive the world around them through learning patterns of tactile stimulation on their torso or tongue over a long period of time that eventually they don't feel the vibrations any more on their body but focus on perceiving the world of objects out there.

One of the earliest forms of sensory substitution was the work of Bach-y-Rita et al. (1969). They built a tactile chair for blind people which converted live visual images from a camera into a tactile display felt on the back of the person sitting in the chair (Figure 2-1). They studied whether this set-up could enable blind people to “see” through the tactile patterns experienced over time on their back. Six blind participants took part. One participant had over 150 hours practice using it and the others had between 20 and 40 hours practice each. During this time they were able to learn to recognise common shapes and objects in their environment and understand how they changed with perspective and when one was placed in front of each other. After many hours of practice, they began to talk spatially about the objects in front of them in their environment, rather than describing the sensations they felt on their back. This suggests that they were literally ‘seeing’ through the vibrotactile technology. Bach-y-Rita & Kercel (2003) proposed the concept of brain plasticity, to explain how the signals go from being a sensation on a person’s back to becoming a perception of the world around them. With repeated practice the brain reorganises itself to process the tactile input in the same way a sighted person might process visual information.

Changes in brain organisation may take many hours of practice, but recent studies have also found that people can learn to do simple tasks using a tactile visual sensory substitution system very quickly. In a recent study (Bird, Marshall, & Rogers, 2009) visual input was mapped onto an array of vibrators placed on the stomach. Blindfolded participants wore a brightly coloured glove and tried to catch a brightly coloured ball. The movement of these two objects were continuously captured by a video camera and mapped onto the vibrators as two moving vibrations. The study showed that participants quickly learnt how to understand this sort of mapping and most were able to catch a ball rolling across a table after a few attempts. It was particularly important in this case that the participants could feel the location of their hand in the display and by moving their hand could affect what they felt on the vibrotactile display, effectively giving tightly coupled real-time feedback in response to movement. Another point to note with this study is that the participants could see the set-up before taking part so they had a visual map in their mind of where they were sitting in relation to the table which they could then project the movement of the ball and their hand on to. The findings suggest that it is the real-time feedback in response to body movement combined with this mental map of the table that enabled participants to learn to catch the ball so quickly. This implies that when designing real-time feedback for motor learning applications, giving participants visualisations or metaphors for mentally mapping the vibrations that they feel on their body to their environment could be one way of helping participants to learn how to respond to them. For example, Lieberman and Breazeal (2007) used the concept of a “force-field” around the correct movement to explain how their feedback for motor learning worked. This enables participants to visualise that there are regions of space which cause vibration and regions which do not.

Figure 2-1: Bach-y-Rita et al.'s tactile chair connected to a camera to convert live visual images into tactile images (Bach-y-Rita, Collins, Saunders, White, & Scadden, 1969). Permission to reproduce this figure has been granted by Nature Publishing Group.

The research on sensory substitution shows that providing training using vibrotactile feedback that maps onto objects perceived in the environment is able to change from an experience of tactile sensation on the skin to a spatial experience of objects in the environment. The evidence suggests that when this happens, the information from the tactile display is no longer being processed consciously by the user. In terms of giving feedback to musicians, being able to design a system that they do not have to consciously process has obvious advantages because music playing is already very cognitively demanding. However, the evidence from these studies also shows that the amount of training time needed to achieve subcognitive processing is lengthy and may not be worthwhile for a musician who wants to experience more immediate effects on their playing that they can understand is helping to improve it. Therefore, although a designer might hope that learners will gradually learn to respond more quickly and intuitively to feedback over time, the main aim of a feedback design should be to help learners to consciously improve their playing.

2.6.2 Feedback for Sensory Augmentation and Navigation

Research has also taken the concept of brain plasticity a step further to investigate how a new type of sense might be perceived. Nagel et al. (2005) developed a belt that used vibrotactile feedback to continuously indicate the direction of North to the wearer. Four participants wore the belt during waking hours for a six-week training period. Before and after this training period the participants were tested in navigations task with, and without, the belt. Their performance in blindfolded navigation tasks while using the belt improved significantly after the training period. However, in a virtual reality navigation task where participants also had visual landmarks to guide them there was very little evidence that training with the belt improved performance. Nagel et al. suggest that this is because the virtual reality environment is not similar enough to the real-world environment where the training took place.

Other vibrotactile belts have been designed to help people to navigate a specific route. Van Erp et al. (2005) built a belt of eight vibration motors evenly spaced around the waist. This was used to indicate the direction of the next waypoint along a planned route. Participants used it to successfully to navigate the planned route between different waypoints which were laid out on an open field but not visible to them. They also found that more complex displays, which used variation in the timing and intensity of the vibrations to indicate distance to the next waypoint, tended to make participants walk more slowly than the simple display and were no more effective for navigation.

Pielot and Boll (2010) tested a similar vibrotactile belt which used vibrations on the waist to indicate the direction of the next two waypoints. They studied it with pedestrians on the busy streets of a town and compared it to using a more traditional visual GPS navigation device which provided maps and arrows. They found that the vibrotactile directions were more practical in a busy street because they did not require the user's visual attention, but that people navigated more accurately using the visual navigation device. Significantly less near accidents occurred with the tactile wayfinder than with the visual device, but there were significantly more errors while using the tactile wayfinder. The errors with the vibrotactile wayfinder usually occurred when there was the choice of two routes heading in a similar direction. This shows a trade-off in the choice of modality: a visual display can give more information but prevents people being aware of the world around them; a vibrotactile display can only give limited information but allows visual attention to be directed elsewhere.

Like Nagel et al. (2005), Pielot and Boll (2010) also report that participants found the vibrations on their waist uncomfortable. They report individual differences in participants' experience of the vibrotactile feedback: half of the participants said they found it tiring to pay attention to the vibrations; the other half said that the vibrotactile allowed them to pay more attention to the environment. Along with the evidence from Nagel et al., this suggests that some people are more in tune with vibrotactile feedback than others.

A contrasting approach to sensory augmentation is “The Question” – a form of immersive theatre experience (van der Linden et al., 2011). It was devised by a group of blind and non-blind actors and researchers around the theme of sensory experience and exploration. Audience members had to navigate their way around the installation in absolute darkness. In the installation there were several zones where tactile stage sets had been created for audience members to explore through touch and where dialogue and sound effects were given through the headphones that each participant wore. Each audience member was given a “Haptic Lotus” to guide them around the performance, this was a flower which opened and closed to indicate to participants whether they were close or far away from a zone of the performance. Sighted participants found that being deprived of visual stimuli was both disconcerting and heightened their other senses, changing the way they experienced and explored the world around them. Blind participants found that being in an experience that centred around touch was liberating because it gave them licence to use their tactile sense in a way that they would like to but is not the socially accepted by sighted people. Participants found the Haptic Lotus to be a reassurance in the darkness and likened the way it moved and communicated with them to it being alive. Despite this, many participants did not believe that they used it to navigate the space; however, observations suggest that participants did use it to scan the space around them and find their way. The findings from this study show that actively depriving the use of one sense can change the way people experience their other senses. This suggests that sensory modalities do interact with one another and that vision tends to dominate our sensory experience. Again this somewhat challenges theories such as working memory (Baddeley, 1992) and multiple resource theory (Wickens, 1980, 2008) that sensory channels run in parallel (see Chapter 3 for detailed explanations of these theories). This study also shows how the presence of real-time feedback that communicates with users can be reassuring, this is similar to the way the application of real-time feedback for physical therapy for chronic pain suffers also reassured participants that they were carrying out safe physical activity.

2.6.3 Feedback for Text Entry

Another area where real-time vibrotactile feedback has been utilised is for text entry in touch screen mobile phones (Brewster, Chohan, & Brown, 2007). Brewster et al. investigated whether *"tactons"* (tactile icons) can be used to improve the accuracy of text entry on a mobile touch screen by indicating to users when they had made a mistake. In this study the tactons used rough or smooth vibrations to communicate different messages. Rough vibrations are when the waveform of the vibrotactile feedback is not a smooth sine wave, but has been changed using amplitude modulation to give a perception of roughness (Brown, 2007). The smooth vibrations were used to indicate when a key had been pressed successfully, the rough vibrations indicated an error in the key press. The vibrotactile feedback is felt on the hands which are much more sensitive to tactile stimuli than the rest of the body (Brown, 2007). This set-up was compared with a normal touch screen. It was found that significantly more errors were corrected in the real-time vibrotactile feedback condition and participants experienced a significant reduction in perceived workload in the tactile condition as measured through a TLX questionnaire. Further research has gone on to confirm that vibrotactile feedback improves text entry on a touchscreen and has shown that with a more complex vibrotactile display performance on a touch screen can be brought very close to that of using a physical keyboard (Hoggan, Brewster, & Johnston, 2008).

This research shows that real-time vibrotactile feedback can encode information in other parameters as well as position such as roughness. However, the vibrotactile feedback in these studies was given on the hand and fingertips which are more sensitive than the body, but are not available to a violin learning application because the fingers are in constant use and the feel of the finger against the string is important to playing. Research has shown that roughness can also be perceived on the arms (Brown, 2007). This study also shows the value of providing real-time feedback about errors in situations where no feedback exists currently. It positively shows that having errors brought to a person's attention in a timely manner enables them to correct them before moving on. Ideally in a violin learning application, this could also become the case, so that learners are made aware of mistake as they practise in a way that they would not otherwise have the ability to notice.

2.7 Summary

The research reviewed here shows how using real-time feedback in human activity can play many different roles: guiding movements; alerting people to mistakes; motivating people; and augmenting or replacing senses. When designing real-time feedback for music learning the aim is for the feedback to help players to improve their technique by improving their movement and posture. Therefore, real-time feedback should either guide movement or alert learners to any mistakes they are making. Alerting learners to mistakes may be the more suitable role, as playing a musical instrument is physically and cognitively demanding and using real-time feedback to guide movement seems to work best in laboratory settings where participants can focus on it fully, whereas alerts to mistakes appear to work better in more demanding situations.

Music learning, like rehabilitation, is also a long-term process and learners need to stay motivated. Therefore, a secondary role for the feedback should be to motivate learners. As Singh et al. (2014) show, feedback which allows people to measure their progress and achievement can be particularly motiving and may be well suited to music learning. The literature about sensory substitution shows that training with real-time feedback has the potential to change the way people use their brain and the feedback may eventually be processed without consciously thinking about it. However, the training time needed to do this would be too long to be motivating or practical for those learning an instrument. However, it is still worth considering how people's use of real-time feedback will change as they get more familiar with it and start to respond to the feedback more automatically.

This review also shows that different modalities are suited to different situations. Vibrotactile feedback appears to be good at getting people's attention, but can sometimes be felt to be too forceful and disturbing. Focal visual feedback can communicate a lot of information quickly and effectively but is already used by violinists for reading music. Ambient visual feedback draws attention less strongly, but this lets people choose how to respond which may make it more enjoyable and motivating to use. Auditory feedback has also been shown to be effective but may interrupt how participants listen to their instrument. Of these modalities vibrotactile feedback or ambient visual feedback appear to be most suitable for learning the violin. These two modalities contrast strongly with one another, but it will require user studies in real learning and practice situations to judge which is more appropriate.

The literature reviewed also reveals questions of how much information should be communicated in real time and how this should be done. Evidence suggests that when applications for motor learning are taken into realistic situations with a lot of external demands the amount of feedback that people can perceive and respond to is greatly reduced (e.g. (Spelmezan, 2012)). Therefore, it is important to be conscious of the information load that a real-time feedback design has and try to find ways of reducing it.

In sum, this review suggests that vibrotactile or ambient visual feedback would be best suited to helping people learn the violin. The feedback should be designed to be used consciously and will need to mediate attention. It should also aim to help learners stay motivated. The technology interventions described here give some guidance on how to design feedback with these aims in mind, but they give little theoretical basis for understanding how feedback might work. Therefore, the next chapter of the literature review brings in theory and experimental evidence from other domains such as psychology and sports science and aims to gain a better understanding of attention, cognition, motor learning and motivation, with a specific emphasis on vibrotactile and visual feedback.

Chapter 3 – Attention, Modality and Motivation

3.1 Introduction

This chapter aims to provide a theoretical understanding of attention, modality, perception and action, and motivation to inform the design of real-time feedback for learning the violin. The goal of real-time feedback in this thesis is to aid practice by enabling players to become more aware of their body in relation to a particular aspect of playing technique they are trying to improve. This could take the form of pointing out mistakes or bad habits which players find difficult to be aware of while playing. Alternatively, it could reward good playing to motivate players to focus on improving a particular skill. Or it could act as a gauge that provides the extra piece of information that makes it possible for a player to understand the movement that they need to make. This is analogous to the yaw string which is attached to the windscreen of a glider so that pilots can see whether they are moving straight through the air (string vertical) or whether they are sliding sideways (string diagonal). Without the hair it would be very difficult for a novice pilot to be aware of this, therefore they focus on it strongly. Over time, they will become more familiar with other cues such as vibrations of the glider which would indicate whether they are sliding and will have to focus less strongly on the feedback from the yaw string apart from at key moments. However, unlike the yaw string, our aim for real-time feedback for learning the violin is that eventually the player should be able to perform the technique correctly without the aid of the feedback.

Based on the review of the literature about technology feedback, my position is that real-time feedback should be designed for learners to consciously engage with it to understand and improve their movement and posture. A key way in which real-time feedback can encourage improvement in movement is its ability to draw attention to things about the way a learner plays (for example mistakes) which the learner would otherwise be unaware of. In such a model of learning with real-time feedback attention and cognitive processing are important for use.

Real-time feedback may also influence learners' playing without them consciously attending to it in the same way we use our ambient sensory system to negotiate our surroundings. However, being able to respond to feedback in this way will need to be learnt and as sensory substitution and augmentation studies show (Bach-y-Rita, Collins, Saunders, White, & Scadden, 1969) (Kaczmarek, Webster, Bach-y-Rita, & Tompkins, 1991) this can take many tens of hours of practice to achieve precognitive reactions to a new 'sense'. Considering that learning to master an instrument requires approximately 10,000 practice hours (Ericsson, Krampe, & Tesch-Römer, 1993), it is not reasonable to expect learners to also spend many practice hours learning how to use the feedback system as well, particularly because new feedback may have to be learnt for each aspect of playing the violin. As such this thesis focuses on designing and studying real-time feedback for conscious use and engagement, if a more automated use then follows after prolonged use this will be an interesting finding but should not be relied upon.

Playing the violin involves the use of multiple modalities; in particular attention to auditory and visual channels is essential for tuning and reading music respectively. In addition players may use the tactile modality to feel vibrations from the instrument, proprioception to place the fingers correctly on the fingerboard and peripheral vision to take cues from other players or watch the conductor. Not only is there high sensory load but also high cognitive load as players translate written music into movement and sound or try to play a piece from memory. In order to do this difficult task, players rely on automated skills which have been proceduralised through practice.

To design real-time feedback to be used whilst playing the violin we need to understand how attention can be distributed between the different activities and how this is affected by factors such as cognitive load. We should also consider how different ways of directing attention affects how we learn proceduralised skills. Moreover as violin playing is a multimodal activity we need to be aware of the properties of different modalities and how they interact with one another.

As well as catching attention real-time feedback also needs to communicate action quickly. Therefore, this chapter reviews literature about how perception and action are linked and whether the design of the real-time feedback stimulus might be used to suggest a particular action to participants. In addition, learning the violin is a long term goal which requires sustained motivation and another possible role for real-time feedback is to help students to stay motivated when practicing. Therefore, this chapter finishes by looking at theories of motivation and how they have been applied to the field of music learning. In order to cover all these areas and answer the questions above this chapter draws on relevant literature from, experimental, theoretical and applied psychology as well as literature from other domains such as education and sports science. At the end of each subsection the evidence is interpreted in terms of implications for the design of real-time feedback for violin playing.

3.2 Attention

The ability for real-time feedback to draw attention to mistakes has a lot of potential to help learners to correct their movement or posture when playing the violin. However in order to design for this we need to understand how attention is directed and mediated and how this relates to learning a skill.

Attention can be categorised into two types: endogenously controlled attention, which is voluntarily directed towards a target; and exogenously cued attention in which something in the environment causes a person to involuntarily direct their attention towards it (Posner, 1980) (Mulckhuyse & Theeuwes, 2010). Exogenous visual cues can take the form of a "singleton" which is something stands apart from its surroundings because it of a difference in colour, form, movement or other property for example our eye is drawn to a poppy in a field of wheat because it is different from the surrounding scene. A particularly strong form of exogenous cue is when something suddenly appears or disappears against an unchanging background, for example a flashing light is particularly good at attracting attention. Exogenous cues such as these "abrupt onsets" have the ability to attract visual attention without the person becoming consciously aware that their visual attention has been diverted (Mulckhuyse & Theeuwes, 2010).

Exogenous cues are particularly relevant to the design of feedback because feedback will need to be able to mediate attention towards aspects of playing when appropriate. To play this role it needs to be able to attract attention without requiring conscious monitoring.

3.2.1 Attention and Skill

Attention and skill interact in unusual and unexpected ways. For example, focusing on carrying out a skill can improve a performance in novices but impair performance in experts. Research from sports science and applied psychology shows that exactly how and where attention is directed changes performance and learning outcomes.

Skill-Focused Attention

Most people may be able to think of times where they have made mistakes in familiar activities because they have attended to them more strongly than usual. Imagine trying to tie your shoes laces while thinking about what you are doing with your hands and where each lace should go – such attention makes the task more difficult. This phenomenon has been demonstrated experimentally. For example Beilock et al. (2002) tested expert golf and soccer players and found that they performed better in a dual task condition (with an auditory distractor task) than when they were instructed to focus on elements of the skill they were executing (the golfers were told to concentrate on the movement of the their putting swing, soccer players were told to attend to the side of the foot in contact with the ball). In a second experiment with the soccer players, novices also performed the same tasks in this case the results were reversed, novices performed better in the skill focused case than in the dual task case. Similarly, if the experts were required to carry out the same tasks with their non-dominant foot, the skill focused case produced the best results.

These experiments show that skilful performance is aided by skill-focused attention when a person is unfamiliar with the skill but as they become more expert attention hinders both accuracy and speed. This is accounted for by the way well practised skills become proceduralised or automated. When something is automated in this way it can be performed faster than doing each step consciously. This procedural knowledge is interrupted if a person then attends in detail to the action they are carrying out.

Four Stages of Learning a Skill

If attention can disrupt practised performance, is it desirable to use real-time feedback to focus learners' attention on their playing in this way? To answer this we must consider that real-time feedback is not intended to improve a single performance, but is intended to aid learning a skill. Here it may be useful to consider the four stage model of learning a new skill (Adams, L. Gordon Training International, 2009). As well as business training, this model has been applied to medical training (Lake & Hamdorf, 2004) and sports coaching. It suggests that as someone learns to become an expert in a particular skill they go through four stages:

1. Unconsciously unskilled: The learner is not able to perform the skill well but is unaware of the mistakes he or she is making or what he or she needs to do to improve.
2. Consciously unskilled: The learner knows what they should be doing but is still unable to perform the skill.
3. Consciously skilled: The learner can now perform the skill but it requires full attention to do it.
4. Unconsciously skilled: The learner can now perform the skill smoothly without attending to each step.

In the case of playing the violin one might progress from stage one to stage two by being taught by their teacher. Moving on from stage two through stages three and four would require practice and repeated playing. Practice will allow aspects of playing to become proceduralised and automated. It is important that the skill is being performed correctly during practice so that the correct version of the skill becomes automated. In other words stages two and three where the learner is aware of what they are doing and comparing it to what they should be doing are vitally important to acquiring expertise. It is in the transition from stage two to stage three that real-time feedback could be particularly useful because it enables players to be aware of their movement and posture and practise them correctly.

Locus of Focus

Another factor in the way attention affects skill performance is where attention is focused. Weiss et al. (2008) and Wulf (2007) have studied the effect on performance of an internal focus versus an external focus when carrying out a motor skill. For example when shooting a basketball at a net an external focus would be focusing on the net, whereas an internal focus would be focusing on the shooting arm. Experimental evidence shows that in general an external locus of focus is more effective for practising and performing motor skills than an internal one for experts as well as learners (Wulf, 2007) even when it is not the person's preferred choice of focus (Weiss, Reber, & Owen, 2008). Moreover retention tests show that learners who learn using an external locus of focus continue to perform better even in trials where their attention is not directed externally by the experimenter.

In one of these experiments (Shea & Wulf, 1999) participants had to stand on a balance board and try to keep it level. On a monitor participants were able to see a representation of the angle of the board; this representation is a form of real-time feedback. One group was told that the feedback represented the angle of their feet (internal locus of focus) and another group were told that it represented the angle of some yellow lines positioned in front of them (external locus of focus). Two other groups had no feedback but were each given either an external or internal focus. Participants performed the task for two days and then carried out retention tests on the third day. Of the four groups those with feedback performed better than those without and the group with feedback and an external focus performed best of all both in the practice sessions with the feedback and in the retention tests afterwards.

Design Implications – Attention and Skill

Learning a complex skill like violin playing requires learning different aspects of playing at different times. Each time learners may go through a cycle of becoming aware how they need to be playing and then practising it until it becomes automated. While working on one aspect of playing they will also be relying on automated processes they have learnt previously. As the research on attention and skill shows, drawing attention to an automated process can damage performance. Therefore it is important that real-time feedback is only used to mediate attention towards aspects of technique that are in the stage where they are being consciously learnt (stages 2 and 3 in the four stage model) without disturbing the automated processes involved in other elements of playing. The choice of feedback could be made under the guidance of a teacher, or through the learner analysing themselves which could be supported by sensing technology to help learners understand their strengths and weaknesses. Alternatively the feedback technology could aim to monitor the performance of the learner and automatically remove feedback when it infers that expert performance has been reached.

When designing feedback it is also necessary to consider where the feedback directs attention - whether it promotes an internal or external locus of focus. As mentioned above, an external locus of focus is often more desirable than an internal one. The feedback itself is an external stimulus but it is linked to movement of the body, therefore it is difficult from the outset to say whether real-time feedback will promote an internal or external focus. It will be down to the details of the design and how the design is then explained to participants as to whether the feedback will promote an internal or external locus of focus.

3.2.2 Attention when Multitasking

Playing the violin requires the player to attend to many things at once for example reading the music, listening to tuning and, in the case of learners, visually checking finger and bow position. In addition, the feedback will also require some of the player's attention. There are several theories, backed up by experimental evidence from dual task experiments, that aim to describe how people mediate attention when multitasking. The theories reviewed here are Working Memory, Multiple Resource Theory and Threaded Cognition.

Working Memory

The early theorising of working memory was that it comprised three parts: a central executive which controls how resources are dedicated and where attention is placed, the visuospatial sketchpad which holds visual and spatial information and the phonological loop which holds speech (Baddeley, 1992). Results from dual-task experiments show that if the two tasks require the same element of working memory then they will interfere with one another.

The idea of dual-processing in working memory has been utilised in multimedia learning (Mayer & Moreno, 1998) where it was found that explaining things using animation and narration was more effective at achieving learning outcomes than using animation and text. Mayer and Moreno (Mayer & Moreno, 1998) explain their findings through the idea that visual and auditory channels are processed separately (at least at a low level) which makes it easier to split attention between visual and auditory information than to attend to two pieces of visual information at the same time. It is important to note that this version of working memory has a slightly different emphasis to Baddeley's (1992). In his version of working memory the verbal information, whether delivered through text or spoken word, would draw mainly on the phonological loop as it is processed through "subvocalization" although the text would also draw on the visuospatial sketchpad as well. Mayer and Moreno's version of working memory places a strong emphasis on the difference in modality by which information is delivered facilitating a "split-attention effect". Whereas in Baddeley's original version of the visuospatial sketchpad and phonological loop it is the nature of the information being processed which is most important.

Working memory load can have other effects than simply interference between two visual or two verbal/auditory tasks. For example, high working memory load, placing load on the central executive, can disrupt the way attention is focused (Lavie, 2005). De Fockert et al. (2001) were able to show this in an experiment where participants had to remember a string of digits which were either high memory load because they changed each trial or low memory load because they stayed the same, whilst doing this they also had to classify written names into two categories (pop-stars or politicians) behind the written names were faces to distract their visual attention. Through fMRI scans it was possible to show that in cases of high load participants paid more attention to the distractor faces than in cases of low working memory load. This demonstrates that working memory load affects the ability to focus attention.

Design Implications – Working Memory

Violin playing should be an activity which places a high working memory load on the player. Since working memory load can disrupt a person's ability to focus his or her attention where they intend to, feedback will need to take the form of exogenous cues which can attract attention when necessary. Moreover, since participants may have difficulties mediating their attention between the feedback and other elements of playing there is a danger that the feedback may become a distraction if it continues to present information to the player when it is no longer necessary for improving their playing.

In music playing, it is hard to divide processes between those which are visuospatial and those which are phonological since it does not involve language in the normal sense. Moreover, if music is seen as a form of language for musicians (an argument that could only be made for experienced musicians), the spatial movement of the body and the positioning of the notes viewed on the page are so intimately linked to the musical sound that it would be difficult to separate the visuospatial components from the phonological. This makes Baddeley's model of working memory difficult to apply here. It may be more useful to take the perspective used in multimedia learning and consider how information is distributed across different sensory channels. The idea being that it would be best to deliver feedback along the sensory channel which is under the smallest load from playing the violin.

Wickens Multiple Resource Theory

Wickens multiple resource theory (Wickens, 1980, 2008) places emphasis on both the importance of modality and type of information (verbal versus spatial that he calls coding) in attention and mental workload. This model of mental resources has 3+1 dimensions:

1. **Stages of processing:** this is split into perception, cognition and reaction
2. **Codes of processing:** this is split into verbal or spatial

3. **Modalities:** this can be visual or auditory and is described as only concerning the perception stage. It has been suggested that other senses may be added here, for example the tactile modality (for example (Boles, Bursk, Phillips, & Perdelwitz, 2007) included tactile workload in their multiple resources questionnaire).

+1. The last dimension is nested within the visual modality and describes two visual channels: **the focal and the ambient visual fields**.

Wickens argues that the more two tasks require different resources the more time-sharing is possible between the two activities and the more efficiently people will be able to conduct the two concurrently. The theory is supported by fMRI results showing that the different resources correspond to different parts of the brain. It is also supported by results from many dual task experiments (for a review see (Wickens, 2002)). However there are also counter examples which would challenge some of the predictions made by this theory. For example, the Colavita visual dominance effect ((Colavita, 1974) see later in this review) challenges the idea that visual and auditory stimuli will not interfere with one another when being perceived; and experiments demonstrating tactile gating (Chapman, Bushnell, Miron, Duncan, & Lund, 1987) contest the separation between perception and action.

Threaded Cognition

Another theory that employs the concept of multiple resources is threaded cognition (Salvucci & Taatgen, 2008). Threaded cognition is a theory of multitasking in which different tasks take the form of threads which can run concurrently so long as they are not trying to use the same resource at the same time. When a resource is unavailable because it is being used by another thread the thread must wait until it is available and this is the explanation for the way two tasks can interfere with one another.

Threaded cognition distinguishes itself from other theories of multitasking by including rules for the central executive which it claims many other multiple resource models rely on but overlook. By this it means that the way the threads are organised and prioritised follow a simple set of rules which passes the resources between them. These rules are as follows: when a thread needs a resource it acquires it as soon as it is available (it is greedy), once it has finished a process using that resource it then hands it on immediately (it is polite). If two threads need a resource at the same time when it becomes available, it goes to the one that has least recently fired a rule on the procedural resource.

This theory has advantages over Wickens' (1980, 2008) theory in that it can be used to quantitatively model the processes occurring in dual-task experiments and make predictions about the speed people will take to carry out one set of tasks compared to another. This analysis has been carried out on situations such as driving and dialling a mobile phone and the predictions of the model fit the experimental evidence well (Salvucci & Taatgen, 2008).

However, to model tasks in this quantitative way they must be very clearly defined. For example, the study of driving limited the driving task to checking that the vehicle is on course and steering and accelerating accordingly and monitoring the car in front in case of braking. Violin playing is more complex than this, it involves continuously reading music, fingering notes and moving the bow to correspond to this, being aware of tuning, listening and responding to the music. With the addition of real-time feedback participants would also be checking the feedback and then adjusting their playing in response to this. In addition there may be unforeseen demands on resources which come from the environment. This would form a complex model of threads which would be difficult to define without a great deal of research into how these processes individually take place. Simplifying the violin playing task would be one way of making this complexity more manageable, however this would sacrifice naturalism. The research in this thesis aims to produce findings which are directly applicable to understanding and designing real-time feedback for violin playing in the real-world; therefore we have chosen not to take this approach.

Design Implications – Multiple Resource Theories

These theories suggest the existence of many independent mental resources which are employed by people when multitasking. If tasks draw on different resources or the use of these resources can be interleaved then these theories predict that concurrent tasks will interfere minimally with one another. If both tasks draw on the same resources then tasks will interfere more strongly. Both theories suggest that if real-time feedback is delivered in a different modality to those used by violin playing it should be easier to perceive than if it uses a modality which is already busy. In the case of violin playing the focal visual modality is used for reading music and the auditory modality is used for listening to what is being played (to check tuning etc.). Also important is the proprioceptive/kinaesthetic sense, which is used to guide finger positioning and bowing movement. The modalities which are used less are the ambient visual modality and the tactile modality. The ambient visual modality would be in use if playing in an ensemble to take in cues from other players and the tactile modality is used somewhat for feeling the vibrations from the instrument and the resistive pressure from the bow. However, compared to the other modalities, these two are used the least, making them the best candidates for giving feedback. This conclusion follows from these theories which claim that different modalities draw on different mental resources. However, there is evidence of cases where one modality can dominate another which would prevent multitasking from being achievable. Therefore, it is important to understand different modalities and the way they interact with one another before fixing on the best modality for giving feedback.

3.3 Modality

Not all modalities are equal. Different modalities have different properties, for example the auditory modality is more useful when temporal accuracy is required (Repp & Penel, 2002) and vision is more useful when spatial resolution is needed (Ernst & Banks, 2002) and touch may be more suited to finding out properties of an object such as texture and hardness (Ernst & Banks, 2002). Modalities also behave differently when combined with one another. For example, reaction times change significantly when people are required to split their attention between more than one modality (Hanson, Whitaker, & Heron, 2009) and in some cases one modality may dominate another (Colavita, 1974).

The following section looks at some of the phenomena which have been discovered when comparing different modalities. Studies which include the tactile or haptic modalities or ambient vision are given precedence in this review because these are the two main candidates for giving feedback.

3.3.1 Ambient versus Focal Vision

Leibowitz et al. (1983) characterise focal or foveal vision as being associated with conscious awareness and used for tasks such as identifying objects and studying their properties. On the other hand, ambient vision is described as part of an ambient system which includes other senses such as balance and can operate without attention. The ambient system is most often used for understanding how the individual fits into their surroundings for spatial orientation, postural control and locomotion (Previc, 1998). It relies on information from a much wider field of vision (180 degrees according to (Previc, 1998)) and has much lower spatial frequency making it suited to tasks such as monitoring slant or ambient motion but not for object recognition.

Based on evidence from studies with pilots and drivers Horrey and Wickens (2004) found carrying out one task with focal vision and one with ambient vision is more efficient than two tasks with ambient vision. However, the best performance was found with two tasks involving focal vision. It could be suggested that although resources may be better distributed in the ambient and focal vision case, focal vision may be better suited to many of the tasks that were studied which made it more effective. Activities which were better suited to ambient resources such as motion judgements worked better as peripheral tasks than those which were more suited to focal attention such as object recognition. Horrey and Wickens also suggest that ambient vision is something which can be developed and improved through practice, giving examples from studies which compare expert and novice drivers. However, they also demonstrate that there are limitations on how much ambient vision can be used for even with practice. For example, hazard perception cannot be effectively conducted with ambient vision alone.

Lamble et al. (1999) showed that peripheral vision operates more accurately the closer it is to the line of foveal vision. In their study, participants had to focus on an LED in-car display whilst driving a car. They were told to brake as soon as they became aware that the car ahead was slowing down. Time to collision was then calculated for each trial (a driving instructor was also in the car to prevent accidents). In each trial the display was positioned in different places which would change the distance between the participant's visual focus and car ahead in their peripheral vision. This distance was found to have a significant effect on the time-to-collision so that when the in-car display was positioned further away from the view of the car ahead the time-to-collision was reduced (i.e. they were slower to respond).

The rest of the research described in this section deals with focal vision as this is a more common part of multimodal experiments. This is relevant in a different way as focal visual load in the form of music reading is a common aspect of violin playing. However, given the differences between focal and ambient vision it would not be fair to assume that results found for focal vision such as visual dominance would be true for ambient vision.

3.3.2 Tactile Perception

Tactile perception studies show that people can accurately distinguish between pressure sensations approximately 4cm apart on the arms and 3.5cm apart on the torso ((Brown, 2007) quoting from (Goldstein, 1999) and (Weinstein, 1968)). These are the two most likely areas of the body where real-time feedback for violin playing might be given, as the hands and finger will need to be kept unhindered for playing. There is not the same data available for vibrotactile stimuli but it is predicted to be lower resolution than pressure stimuli because the receptors that sense vibration (Pacinian corpuscles) are large with undefined borders whereas those that sense mechanical pressure are smaller (Brown, 2007), (Gallace & Spence, 2014)). Experiments have also shown that spatial acuity also increases if the position of the vibration is near to an anatomical reference point such as the wrist or elbow on the arms (Cholewiak & Collins, 2003) or the spine or navel on the torso (Cholewiak, Brill, & Schwab, 2004).

3.3.3 Reaction Times to Different Modalities

Different modalities have different reaction times associated with them. In unimodal speeded reaction trials people reacted to auditory stimuli quickest (161.3ms), then tactile (192.3ms) and then visual slowest of all (206.9ms) (Hanson, Whitaker, & Heron, 2009). However, in bimodal or trimodal trials where participants had to attend to more than one modality, results are quite different. In these trials, reaction times to auditory or visual stimuli were significantly increased whereas the response time to tactile stimuli remained almost the same (Hanson, Whitaker, & Heron, 2009). As a result in both trimodal and bimodal cases, response times to tactile stimuli were lower than those to auditory or visual stimuli. In the multimodal trials participants had to respond as quickly as possible to a stimulus from a single modality and were only required to indicate its presence not its modality. A different effect occurs when participants also have to discriminate which modality is present this is known as the Colavita effect.

3.3.4 Colavita Visual Dominance Effect

The Colavita effect (Colavita, 1974) is an effect in which the visual modality dominates other modalities during speeded discrimination tasks. It was originally demonstrated comparing the visual and auditory modalities. In this experiment participants had to press different buttons in response to a visual signal or an auditory signal. However, mixed into these trials, there were also cases where the visual and auditory signal would come on at the same time. The results of this experiment show that the visual modality clearly dominated the auditory modality to an extent that in some case the auditory signal was not even perceived. Since this first experiment many more experiments have been conducted (e.g. (Colavita, 1979) (Koppen & Spence, 2007a) (Koppen & Spence, 2007b)) in which participants were told that there may be some bimodal trials, in all these cases visual dominance has been found, although to a smaller extent than in the original study.

The Colavita effect is not confined to audio-visual studies; it has also been shown in tactile/haptic-visual experiments (e.g. (Hartcher-O'Brien, Gallace, Krings, Koppen, & Spence, 2008), (Hecht & Reiner, 2009)). Hartcher-O’Brien et al. conducted an experiment where participants had to differentiate between a vibration on the finger and a visual stimulus, participants then had to indicate whether the stimulus was visual or tactile by releasing pedals with their feet. They found a similar significant visual dominance effect. Hecht and Reiner (Hecht & Reiner, 2009) compared three modalities visual, auditory and haptic. They presented the stimuli unimodally, bimodally and trimodally. Participants had to indicate by pressing buttons which modalities were present. They found significant visual dominance in the audio-visual and haptic-visual bimodal trials; in both cases participants were more likely to mistake the bimodal case as being visual only than either haptic or auditory only.

The Colavita effect shows that for certain tasks, there are cases where the perception of different modalities may obstruct one another. Playing the violin from sheet music relies heavily on the visual modality and although multiple resource theory and threaded cognition suggest that this should not interfere with the perception of vibrotactile feedback, the Colavita effect challenges this assumption.

3.3.5 Gating of Tactile Perception due to Movement

Movement can prevent tactile stimuli being detected by people. Chapman et al. (1987) found that participants were less likely to detect an electrical tactile stimulus on their arm when actively moving that arm or having the arm moved for them (passive movement) when compared with cases where the arm was stationary. Moving the opposite arm did not have this same effect. Once the stimuli were above a threshold where they would be regularly detected, arm movement was found to have no effect on participants' abilities to discriminate between the intensity of different stimuli when compared with no movement. Nor did it effect participants' subjective perception of the strength of a stimulus.

Post et al. (1994) ran a similar experiment using vibrotactile feedback. They found that a participant's ability to detect a vibrotactile signal was reduced when their arm was moving compared to stationary. In this experiment the vibrotactile stimuli were given using an array of pins which vibrated and the objective measure of intensity was how many pins were set to vibrate. Above the threshold where participants could regularly detect the stimuli Post et al. found that participants' ability to discriminate between different intensities of vibrotactile stimuli was not significantly affected by movement. Movement was found to have an effect on participants' subjective ratings of the intensity of the vibration on a scale of their own choosing. Vibrations were considered to be less intense when the arm was in motion.

More recently Buckingham et al. (2010) have investigated tactile perception in reaching tasks using two hands. Participants were asked to detect which hand they felt a vibration on whilst reaching for points on a screen in front of them. They showed a pronounced and significant drop in their ability to correctly detect the vibration when the onset of vibration was close to the time they were moving their arms for the reach. A follow-up study confirmed that this suppression of tactile sensitivity begins to occur even before the muscles are activated for movement. This evidence is consistent with the findings of another study (Voss, Ingram, P., & Wolpert, 2006) which showed sensory attenuation in response to activating parts of the brain associated with movement preparation.

3.3.6 Individual Differences in Perception of Different Modalities

Another question we might consider is whether all people react to a particular modality in the same way? For example, gender is important when considering subjective sensitivity to vibrotactile stimuli (Neely & Burström, 2006). Women find vibration more uncomfortable and rate the subjective intensity of the vibration higher than men. However the threshold at which they can detect vibrotactile stimuli is the same for both sexes. This threshold does change with age. Sensory thresholds increase with age meaning younger people are able to detect gentler vibrations which older people would not (Bartlett, Stewart, Tamblyn, & Abrahamowicz, 1998).

One currently controversial (but also popular) theory in education suggests that individual learners have preferences to learn using particular modalities. The theory claims that learners fall into three or four different categories. These are: Visual learners, Aural learners, Kinaesthetic learners and learners who prefer reading/writing (Fleming, 2006). There is still a lot of debate about whether learning styles are a useful concept for teaching and learning (Sharpe, Bowker, & Byrne, 2008) or have any physical basis in the brain (Geake, 2008). However the general idea that people have individual preferences towards using particular modalities should not be overlooked. There does not seem to be any clear evidence to suggest any divisions along gender lines (e.g. (Slater, Lujan, & DiCarlo, 2007)).

3.3.7 Design Implications – Modalities

Multiple resource theory suggests that tactile or ambient visual modalities are suited for real-time feedback because they are the modalities used least in violin playing. Ambient vision, however, has limitations. The ambient visual field is more suited to detecting movement rather than tasks that require more conscious awareness like object recognition. This would suggest that an ambient visual display would need to use movement or change to communicate information to learners rather than form or shape. The size and position of an ambient visual display is also important. Ambient vision has low spatial frequency which means that the greater the proportion of peripheral vision the display takes up the easier it will be to use with ambient vision (this is different to just making a large display as the amount of the player's view it fills is also dependent on its distance from them). Secondly, the display should be easier to use if it is positioned close to written music, as ambient vision operates more effectively closer to the visual focus.

Ambient activities often happen in the absence of conscious awareness. However, our model of real-time feedback requires conscious attention some of the time in order for a learner to interpret it and adjust how they play. Therefore, using only ambient vision may not be the most efficient way for feedback to work. It is more likely an exogenous cue such as movement in the ambient visual field will briefly attract focal visual attention which will be used for reading and interpreting the display. This may be more desirable since it appears that people perform better in two centrally attended tasks than in one ambient and one central.

Tactile perception is also limited. Studies show that in some cases it can be dominated by vision. However, it is difficult to predict how this will affect the case of feedback for violin playing because vision is being used for reading music which is not a wholly spatial task where vision has been shown to dominate most commonly, nor is it an attention catching visual stimulus of the style used in studies of the Colavita effect. The study of reaction times shows that in terms of speed of catching people's attention, vibrotactile perception does not suffer in the same way as other modalities in a multimodal situation. This may be useful, as bringing attention to a problem in a timely manner may be the key contribution of the feedback to learning the violin. Tactile gating caused by body movement is important to be aware of, as violin playing involves large movements. The tactile stimuli where gating was shown were close to the threshold between where people can feel the stimulus and where they cannot, as long as the stimuli were well above the threshold, detection of the stimuli was unaffected by movement. This suggests that as long as the feedback stimuli are well above threshold the movement of playing should not interfere with detection of the tactile feedback. However, given that movement has been shown to affect people's subjective experience of the feedback and players are also under cognitive load and sensory load from other modalities, it is difficult to predict whether movement will affect perception of the vibrotactile stimuli.

Tactile gating is particularly interesting because it is a case of action interfering with perception. Wickens multiple resource theory separates action and perception into two sets of resources, however tactile gating would suggest that this distinction is not clear cut. There are other theories such as the Theory of Event Coding (TEC) (Hommel, Müsseler, Aschersleben, & Prinz, 2001) and Embodied Interaction (Dourish, 2001) which suggests a much more intimate link between action and perception.

3.4 Perception and Action

In order for real-time feedback to be effective, a link must be made between perceiving the feedback and taking action to improve playing. This section reviews three different theories which explore the link between perception and action.

3.4.1 Gulf of Execution and Evaluation and 7 stages of User Action

A classic theory for understanding user actions when interacting with technology is Norman's 7 stages of user action (Norman, 1986). He examines how users try to carry out goals using technology but are sometimes thwarted if there is not enough information to help them move between their psychological goals and a change in the system state. He suggests that as users interact with technology they go through seven stages which normally happen in the following order:

- Establishing the Goal
- Forming the Intention
- Specifying Action Sequence
- Executing the Action

- Perceiving the System State
- Interpreting the State
- Evaluating the System State with respect to Goals and Intentions

As users go through these stages they are negotiating gulfs that exist between themselves and the technology (see Figure 3-1). The first is the gulf of execution (Intention, Action Sequence, Execution) in which the user tries to translate their own goals into actions which will cause the technology to produce the desired outcome. The second is the gulf of evaluation (Perception, Interpretation, Evaluation) in which the user tries to take in the system state and judge whether the goal has been achieved. The similarity between the user's mental model of how the technology works and its actual mechanism affects how wide these gulfs are. Normally, a difference in mental models will increase the gulf of execution. The speed at which the user gets feedback from the technology will affect the gulf of evaluation and their ability to understand whether their actions are creating the outcomes they expect.

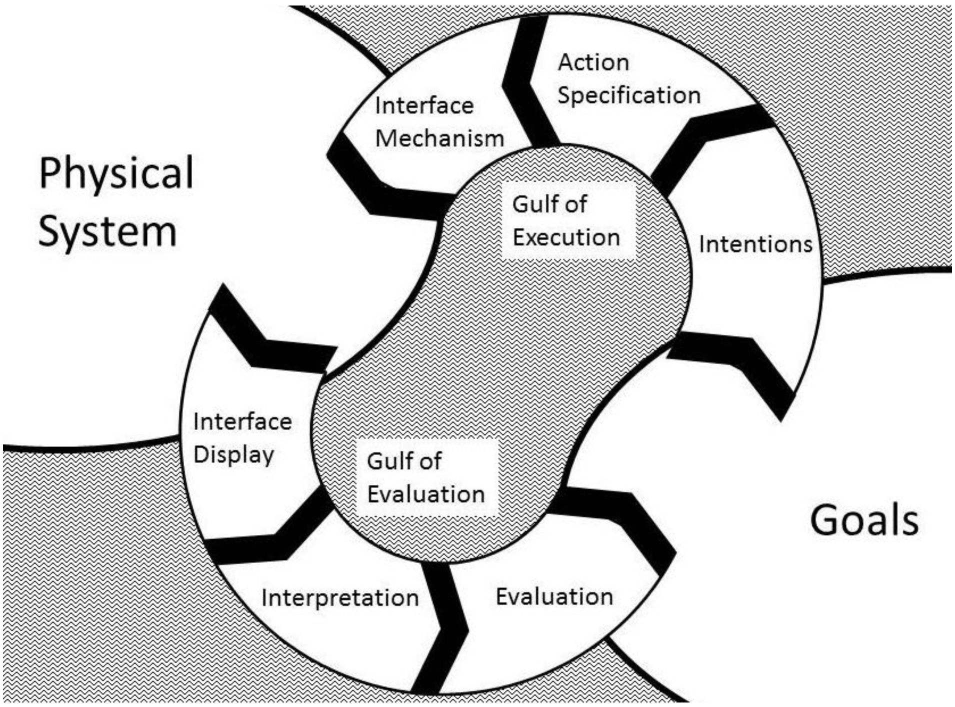

Figure 3-1: How the 7 stages of user action can bridge the gulfs of execution and evaluation based on (Norman, 1986).

Using this model to analyse real-time feedback for playing the violin brings out some interesting design issues. Unlike most applications it is necessary to consider the gulf of evaluation first because real-time feedback comes in response to playing movements rather than in response to the user trying to change the system state. Norman suggests that interpretation of the system state is aided by the speed at which the system gives feedback, as this is almost immediate in real-time feedback the interpretation step will be aided by this. However, before this happens it is necessary to perceive the system state and this requires the system to have caught the player's attention and be perceptually easy to take in the information with only limited attention available. The final stage involves evaluating the system state with respect to the goals of the user. This means that the feedback that the system is giving needs to be relevant to the goals of the users while they are playing the violin e.g. if the user is trying to improve their bowing then the feedback needs to be about bowing.

Once the player has understood what the real-time feedback is communicating about their playing, the gulf of execution must also be bridged. First, by the user forming an intention to correct or improve the aspect of playing they have just received feedback about and then by planning and executing actions to change their playing. How they plan and execute this action will be informed by two factors: the feedback and the player's own knowledge. On the one hand, the feedback itself could suggest action, for example the push-metaphor used by Spelmezan et al. (2009) in their vibrotactile instructions is a cue which is designed to prompt action. On the other hand, feedback could leave it to the player's own knowledge about the movements they are aiming to make to inform how they change and adapt these in response to the real-time feedback. Both approaches have merits. The approach where the feedback suggests action may allow the player to correct themselves more quickly and may be less cognitively demanding; the approach where the player chooses their own response may make the feedback more flexible to different learning goals.

3.4.2 The Theory of Event Coding (TEC)

The theory of event coding (TEC) (Hommel, Müsseler, Aschersleben, & Prinz, 2001) suggests that perceived events (perception) and "intended-to-be-generated events" (action) are both coded in the same way and both stored together in a common representational medium. This would mean that action and perception are almost equivalent in the way they are treated by the brain and there will be interaction between the two. Tactile gating may be thought of as one such interaction. Another example is the way we mainly perceive properties of an object which are relevant to our planned action, such as when looking at a watch people only notice the features of the watch which our relevant to the intended action of telling the time (e.g. angles of the hands with respect to the angle of the watch face) but do not perceive or remember irrelevant information such as the colour of the hands.

If action and perception are coded in the same way then this has implications for the design of real-time feedback. It suggests that the way feedback is presented and perceived by the user will affect how easy it is for them to take action about the feedback. Certain ways of giving feedback that resemble the way that the movements to be made are coded by the brain may be quicker in suggesting the correct action to participants. Similarly, one might expect the intention to take action in response to the feedback to affect how it is perceived. There may also be interaction between the actions involved in playing and the perception of the feedback. One might expect these to be small when the player is mainly using proceduralised knowledge as this does not have the same level of planning which TEC focuses on when discussing action. However, if a player was playing something where they have to consciously plan each note they are playing, TEC would suggest that this will strongly interfere with the ability to perceive real-time feedback. Moreover, if feedback were to be given about two different playing movements, TEC would suggest that the process of taking action in response to the feedback about one of the movements may hinder perception of the feedback about the other type of movement.

3.4.3 Embodied Interaction

Embodied interaction (Dourish, 2001) is a theory used in HCI and the design of interactive systems. It centres around the concept of *Embodiment,* which can have many definitions but at its core concerns being in the world and acting and thinking with or through the body in real time and space. Embodiment is grounded in the philosophical tradition of phenomenology which focuses on human experience as the source for understanding the world. People experience and interact with the world and other people on a daily basis through their bodies in real-time. This embodied experience is the source for understanding everything, even abstract concepts. Rather than delineating between mind and body, embodiment argues that they are the same system. *"The nature of being – how we exist in the world – shapes the way that we understand the world, because our understanding of the world is essentially an understanding of how we* are *in it"* (Dourish, 2001, page 107). Not only that, but objects and our environment can also be incorporated into this system so that a tool that is *"ready-to-hand"* ((Dourish, 2001) quoting Heidegger on page 109) is used as an extension of the body without consciously registering it as a separate entity. For example, in the hands of an experienced musician the bow becomes an extension of the musician's body which is used for creating sound without consciously thinking about how the bow is a separate object to the rest of their body.

With the theory of Embodied Interaction, Dourish is arguing that as designers of interactive systems we should design to appeal to the way people understand the world and create meaning in their everyday lives through embodiment. To do this we should think about the ways that people negotiate their everyday environment through sensing and acting with their bodies; and the way that people understand social situations through interpreting others actions through their own ability to act. By designing in this way we can create systems that can seamlessly become part of people's everyday life by drawing on the skills they already have from their normal experience to understand and create meaning. Moreover, when it comes to evaluating prototypes, we need to use a method that allows participants to act in a natural embodied way and allows them to incorporate the technology into their everyday environment.

In terms of designing real-time feedback this has many implications. Real-time feedback by definition is a process that unfolds in real-time and so lends itself to an embodied perspective. Embodied interaction challenges us to think about how the feedback display can draw on everyday sensory experiences to communicate with participants about their playing. In his design principles Dourish states that *"users, not designers, create and communicate meaning"* (Dourish, 2001, page 173). Not only is this the case, but the technology is not the only factor participants are using to create meaning about their playing; they will also be drawing on their own knowledge and previous experiences and the support of others such as their teacher. However the designer can encourage particular interpretations of the feedback by users which will enable them to take action on the feedback which is helpful to their playing. This is where understanding embodied interaction can play a role. To design in this way will not only require learning about the way people experience feedback through their bodies, but also understanding how environmental factors such as setting, teachers, other players and learning goals influence the way they use the feedback.

3.4.4 Design Implications – Perception and Action

The review of relevant theories described here suggests that we need to design for several stages of interaction when considering how to provide real-time feedback. First, feedback needs to be perceived and then interpreted in relation to the player's goals. Therefore, real-time feedback needs to be giving information that is relevant to the goals of the learner. Next the learner needs to plan and take action. TEC suggests that action and perception are interlinked and therefore it may be possible to design feedback that suggests particular actions to the user. The user's learning goals can also be interpreted as an intention to take action and therefore will also affect how real-time feedback is perceived. Again this highlights the need to design with users' learning goals in mind. However, since there are such a wide range of possible learning goals within the scope of learning the violin this becomes a complex design challenge.

As well as learning goals, embodied interaction brings out many other factors that may influence the way feedback is interpreted and how action is taken. For example, environmental factors such as the setting and the teacher; and factors specific to the way the body is used to play the violin. When people play the violin they interpret that they are making music through their body and the instrument as an extension of their body. When feedback is given while people are playing, the design needs to take into account the perspective of the active player and give feedback that informs or guides movement in a way that is meaningful to the embodied process of playing the violin. As Dourish emphasises, users not designers create meaning and it may be that there needs to be flexibility in the design so that players can appropriate it and make the feedback meaningful for them. In the next section, we consider the role of motivation in playing and how this influences the choice of real-time feedback and when and how to present it.

3.5 Motivation

Motivation can be defined as *"an internal state which serves to activate or energise behaviour and give it direction"* (Huitt, 2001). Some definitions also include the aspect of persistence in behaviour. The motivation to energise a certain action may be quite different to that which will make a person persist in an action. For the purposes of analysing learning the violin, which is a particularly long-term goal, persisting in practising over many years is an important aspect of being a motivated player. There are a number of structures through which motivation can be viewed. Below, a number of these approaches to motivation are described.

Intrinsic versus extrinsic

Intrinsic motivation occurs *"whenever a person experiences himself to be the locus of causality for his own behaviour"* whereas *"when a person perceives the locus of causality for his behaviour to be external to himself"* then he will be extrinsically motivated (deCharms, 1968). Being the *"locus of causality"* means that a person perceives him or herself as the cause or reason for his or her behaviour; if this is external, then the person sees their environment or another person as affecting or controlling their behaviour.

Extrinsic and intrinsic motivational factors do not always interact in an additive way. In certain situations the addition of an external motivational reward can actually act to reduce intrinsic motivation (e.g. (Calder & Staw, 1975); (Greene & Lepper, 1974)). For example Calder and Staw devised two tasks of equal difficulty, one more interesting than the other (i.e. more intrinsically motivating). They found that those doing the more interesting task for money found the task less enjoyable than those doing it for free. Conversely, the monetary reward increased the enjoyment of those doing the less interesting task. This shows how in some cases, a salient, contingent reward can reduce enjoyment of an intrinsically motivating task. However, rewards which are unexpected or which signify competence in a worthwhile task can increase intrinsic motivation (Hallam, 2002).

Attribution theory

Attribution theory (Weiner, 1985) states that motivation is related to our perception of the causes of our successes and failures. These factors can be stable or unstable, controllable or uncontrollable, internal or external (see Table 3-1 for examples). Based upon our attribution of the reason for successes and failures in the past we make predictions for the future. If, for example, a person believes that they failed their exam due to a stable cause out of their control, such as lack of aptitude for the subject or the difficulty of the exam, then the student would not be motivated to study hard to retake the exam. Whereas if the student believes he or she failed due to a controllable cause, such as lack of effort, or an unstable cause, such as bad luck, they will be more motivated to continue their studies. It is generally thought that students are most persistent when they attribute their **successes** to internal controllable factors, such as **effort,** and internal uncontrollable stable factors, such as **aptitude,** but attribute **failure** to unstable factors, particularly those over which they have control, such as **effort**.

Factors	Controllable		Uncontrollable	
	Internal	External	Internal	External
Stable	Learnt ability		Aptitude	Difficulty of specific task
Unstable	Effort	Choice of task	Mood	Luck

Table 3-1: Factors perceived as affecting success\failure

Expectancy x Value theory

A related approach is the expectancy value theory of motivation (e.g. (Eccles & Wigfield, 2002)). This places emphasis upon the individual's expectancy of success or failure in a task but takes into account the individual's perception of the value of the task as well. The main components of an expectancy value theory are:

Expectancy components Students' beliefs about their ability to perform a task. Taking into account their view of their inherent ability, the amount of effort they are willing to put in, time constraints etc.

Value components Students' perception of the value of the task. Task value can be split into four components:

Attainment value. The personal value of doing well on the task. This comes from how well the task fits with an individual's self-concept and sense of identity.

Intrinsic value. The innate interest or pleasure to the individual of the actual activity itself.

Utility value. Whether a task fits in with a person's aims and goals both long term and short term. Other reasons to do the task such as to please a parent or to get a reward.

Cost. Negative aspects of doing a task, e.g. performance anxiety, effort required, risk of failure.

An individual needs to perceive a task as having both an expectancy and a value component in order for the task to be motivating. For example, a student working towards an exam will only be motivated if he believes he has a chance of passing it and that the exam is valuable to him in some way. The value of the exam may come from the sense of achievement the student will get if he passes (attainment value), his intrinsic interest in the subject (intrinsic value) or the fact it progresses his career (utility value). The cost of working towards the exam will be the amount of time and effort he has to expend and the risk of failure.

Learning versus performance related goals

In the study of motivation in education, students' goals are often split into performance or learning related goals. Performance related goals (also known as ego goals) are aimed at getting positive recognition of competence and avoiding being judged as incompetent. Learning (or mastery) related goals aim to understand new things or progress in learning a skill. The focus here is on gaining mastery and knowledge rather than social recognition of ability. Studies have shown that students with learning related goals adopt better learning strategies and apply their learning more effectively (e.g. (Ames & Archer, 1988)); whereas, students with performance related goals who perceive themselves as having low ability are more prone to helpless responses (Elliott & Dweck, 1988). A helpless response occurs when an individual deems him\herself as unable to control a situation and will therefore make no effort to take control even when this is not the case.

3.5.2 Motivation in music

Much research has been carried out to identify factors that influence motivation in music and how these change depending upon age. Zdzinski (1996) found that parental involvement correlated strongly with attitude towards music in high school age children. There was no significant correlation at elementary school; however parental involvement did significantly affect achievement at this age. This suggests that in the early years parents act more as extrinsic motivators but as the child gets older they begin to share the positive attitude towards music displayed by the parent (as implied by (Sichivitsa, 2007)). This correlation could be due to direct parental influence or due the child's higher expectancy of success thanks to early success due to parental involvement. Many other works (e.g. (Lehmann, Sloboda, & Woody, 2007) and (Sichivitsa, 2007)) also place strong emphasis on the influencing effects of parents upon the motivation of their children.

Teachers and conductors have been shown to be influential in motivating musicians of all ages ((Sloboda, 2004), (Galvao, 2002)). In the later years of learning through to professional level, teachers and conductors who are high status role models have been found to be more motivating (Galvao, 2002) whereas in early years teachers viewed as warm and sympathetic were reported to be important (Sloboda, 2004). Sichivitsa (2004) found that the influence of peers had little effect at lower school levels but increased with age. Lehmann et al. (2007) highlight how adolescent players in a music ensemble or band act to indirectly motivate each other since their standing within their social group is strongly connected with their musical ability. However, social groups can also act to discourage more self-conscious musicians if the element of competition is too strong. Gender also plays a role. More girls play instruments than boys and gender affects their choice of instrument (Hallam, Rogers, & Creech, 2008). Despite the fact that boys and girls have performed equally in measures of musical ability, girls generally perform better in music exams. Girls are also more positive in their beliefs about their competence and the value of instrumental music (Hallam, Rogers, & Creech, 2008).

Smith (2005) researched the effect of goals orientation and implicit theories of ability upon music practice strategies among students in music colleges. This study found that learning related goals were closely correlated with an incremental view about ability. Learning related goals led to more varied practice strategies. Students with learning goals were more willing to use non-playing practice methods to improve musicality and expressivity whereas those with ego goals (performance related goals) were negatively correlated with this. Smith also makes a distinction between ego approach goals (the wish to demonstrate high ability compared to others) and ego avoid goals (the wish to avoid demonstrating lack of ability compared to others). There were some findings to suggest that during practice those with ego avoid goals were more sensitive to mistakes leading to more stopping and starting and more time spent retuning their instrument which may not be seen as an effective practice strategy. Table 3-2 summarises the findings of the studies which look at musicians' motivation.

Factor	**Details**	**References**
Parental involvement	Important throughout life although can also have a detrimental effect.	(Zdzinski, 1996), (Sichivitsa, 2007), (Lehmann, Sloboda, & Woody, 2007)
Influence of teacher	Important throughout life. In the early years a sympathetic teacher is best. In a later years of learning an inspiring role model is more motivating.	(Sloboda, 2004), (Galvao, 2002)
Influence of peers	Unimportant at primary school age. Increases with age up to college level. Can have both positive and negative effects on motivation.	(Sichivitsa, 2004), (Sichivitsa, 2007), (Lehmann, Sloboda, & Woody, 2007)
Previous musical experience	Considered important in creating self-concept within music.	(Sichivitsa, 2007)
Perception of ability	An incremental view upon ability is strongly related to learning goals	(Smith, 2005)
Goals	Learning goals lead to a more varied use of practice strategies	(Smith, 2005)
Gender	More girls play instruments than boys. Girls are more positive about their abilities and the value of instrumental music.	(Hallam, Rogers, & Creech, 2008)

Table 3-2: Factors affecting motivation in musicians

Models of motivation in music

"Human beings have a "love affair" with music" (Lehmann, Sloboda, & Woody, 2007) and many people dedicate their lives to mastering the ability to make it. For this reason, music is a particularly interesting area in which to formulate theories and models of motivation. On the one hand, listening to and playing music is seen as intrinsically pleasurable, while on the other, deliberate practice is generally viewed as a boring activity (Galvao, 2002). In the following section three different approaches to motivation in music are examined.

Hallam (2002) has put forward a model to describe how the motivational factors in the literature interact with one another (see Figure 3-2). It is the culmination of an extensive review of studies in the area of musical motivation. She describes how motivation to be involved in music comes from the interaction between an individual's characteristics, self-concept and goals, and the environment in which they are developing. The mediator between these two is the cognitive process of the individual, how he or she interprets the environment in relation to him\herself. Here, she includes ideas from attribution theory showing that how a person interprets the causes of success or failure in the outside world directly affects their self-concept.

This model is a good way of summarising factors that can affect motivation in music. However, it gives no sense of the strength of interaction between these factors. Similarly, everything is shown to interact in a two-way fashion with no sense of what is the predominating direction of dependence. It has not been tested and indeed it would be very difficult to test given the generality of statements such as "the environment". It gives a sense of the complexity of the interaction between the factors which affect motivation but makes no attempts to simplify them into threads of dependencies that would lend themselves to being tested in a traditional manner.

More recently Sichivitsa (2007) has put forward a model with a strong emphasis on expectancy value theory. She proposes that an individual's self-concept in music is dependent upon their previous musical experience and parental support of music. This self-concept in music is linked through academic and social integration to the value placed upon music. She then argues that the value of music will be directly related to the individual's intentions of continuing in music. This model is in part a positive feedback loop since previous experience in music affects the motivation to continue in music.

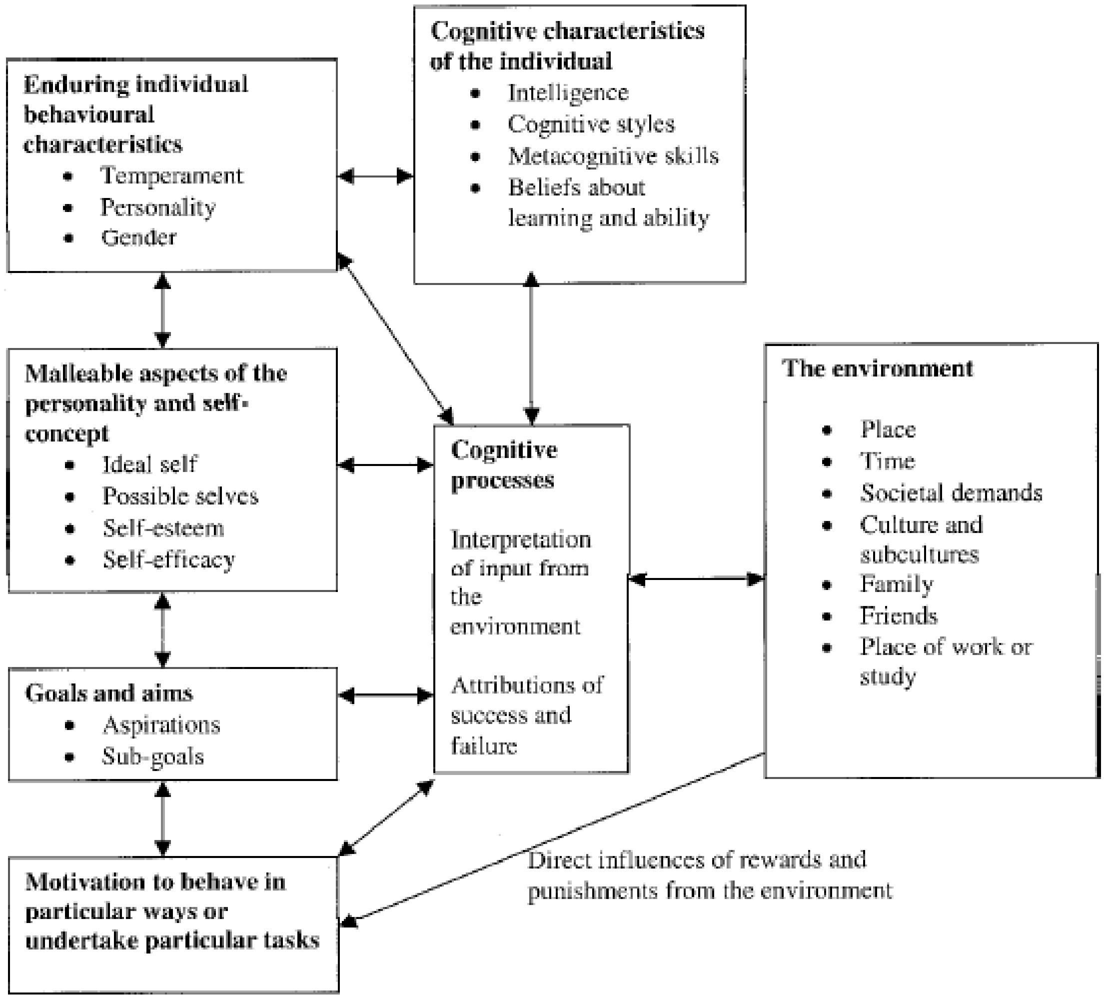

Figure 3-2: Hallam's model of motivation in music (Hallam, 2002). Permission to reproduce this figure has been granted by Taylor and Francis Group.

In contrast to Hallam, Sichivitsa has taken steps to validate it using a questionnaire given to college students (not studying music) who were members of a choir. The questionnaire aimed at measuring each of the factors given in her model and from this data interaction between these factors could be inferred via path analysis. The path analysis indicated a good fit between the theory and the data she collected. However it is not clear whether another model which placed different dependencies between the same set factors would not also be a good fit. Similarly, as acknowledged in her discussion, a correlation between two variables gives no indication of directional cause and effect. For example, parental support of music may grow out of observation of their child's self-concept in music rather than the reverse.

In her discussion, it is clear that she views self-concept in music as similar to students' expectancy of success. However, this term also infers a strong sense of value to the task itself not simply the ability to do it. Equally an individual can place a value upon something without being able to do it, hence the reasoning behind the expectancy-value model treating value and expectancy as separate (but interacting) variables. Treating self-concept in music as expectancy of success would make value dependent upon expectancy. This seems a clear move away from traditional expectancy x value models.

Unlike Hallam's model of musical motivation, Sichivitsa's model only focuses on external influences upon motivation in music. It does not take into account inherent qualities of the individual that may contribute to musical motivation. Taken on its own, the model would not account for the many examples of individuals, such as Handel, who from a very early age pursued the learning of music despite parental disapproval. However, seen as a partial model, it is useful because it places emphasis on key external factors relating to musical motivation, with ideas about how they may interact with one another. Since it is better defined and there are fewer variables involved, it is inherently more measurable and testable than Hallam's.

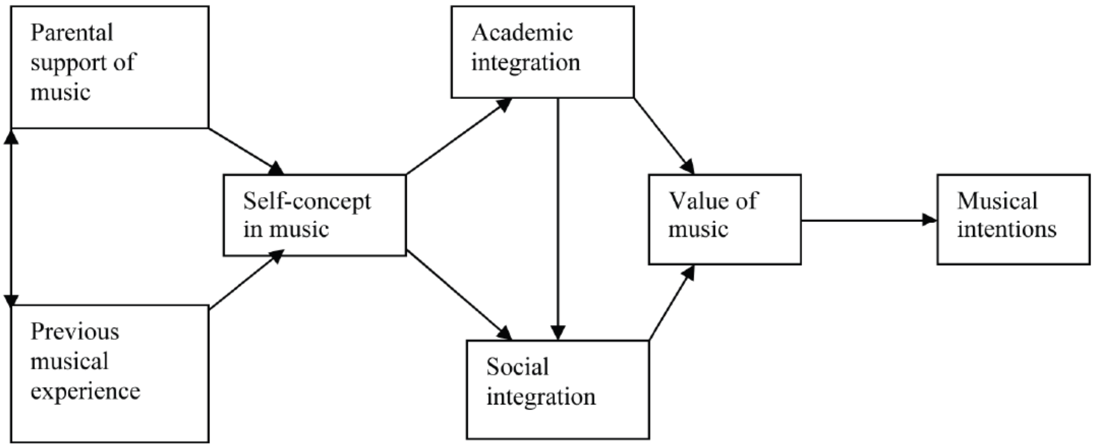

Figure 3-3: Sichivitsa's model of motivation in music (Sichivitsa, 2007). Permission to reproduce this figure has been granted by Sage Publications.

A third approach to motivation in music comes from the book *"Psychology for musicians: understanding and acquiring the skills"* (Lehmann, Sloboda, & Woody, 2007). It is not a model but rather a chapter detailing the complex mixture of factors both external and internal which can affect a musician's motivation throughout his or her life. In the section on intrinsic motivation in music it states that motivation has its roots in pleasurable childhood experiences of music. It then goes on to put forward that a sense of exploration, freedom and personal autonomy can maintain intrinsic motivation in music through to adulthood. The "intensely powerful" experience of musical performance is also seen as a motivating factor, especially for inspiring younger musicians.

The effects of extrinsic motivators are also examined in this chapter, covering similar ground to the review at the start of this section. The effect of a person's beliefs and values upon their motivation is also examined at length. Here, the authors bring in the concept that an individual's ideas about their own ability will affect their motivation. This includes ideas from attribution theory, suggesting that the reasons students give for their successes or failures will affect their motivation. They also note that as music students become more advanced they increasingly take on the view that their successes and failures are due to their 'talent' rather than their effort. The writers ascribe this to the increasing divide which is seen as children enter high school between the high achievers that have *"distinguished themselves"* and the rest of the children studying music. An expectancy-value view on motivation is introduced here, highlighting the many students who continue to learn music not because they think they are good at it, but because they think that music is important and valuable. They go on to note that the four value components outlined in the expectancy-value model described earlier are good predictors of students' future involvement in music.

Based upon this analysis of motivation the book then offers advice to musicians and teachers on how to keep themselves and their students motivated. The main emphasis is upon playfulness and exploration, especially for younger children, but also warning teachers of more advanced pupils to *"guard against taking all the fun out of music!"* It also encourages teachers to help students set appropriate goals, and remind students why they are performing; with an eye to promoting learning and mastery goals rather than performance/ego goals. Adult musicians are encouraged to motivate themselves by *"exerting personal choice"* in some way within their playing.

This description is not intended to be a model; so simplifying the picture of musical motivation is not its aim. Its approach differs from the other two in the way it emphasises the intrinsic aspects of motivation in music and tries to pinpoint some of the places that this comes from and how these can be harnessed to make music learning a more motivating experience. It claims that early childhood experiences will affect motivation into adulthood and that it is important that these are fun and pleasurable. These claims are based upon interviews that showed that most successful musicians had positive memories of music from their early childhood. However, it cannot be certain that these events in their childhood are different from the events in non-musicians' childhoods; for it could be that an inner motivational force meant that the musician-to-be found these activities fun and memorable and the other children without this sense of intrinsic motivation did not. Nor can it be ascertained whether these memories have not been coloured by perspective of their later careers. These could become defining memories because the speaker has such a strong self-concept within music in the present. Nonetheless, it is reasonable to suggest that enjoyable and active musical experiences in early childhood will act to create a lasting view that music is a pleasurable activity.

The ideas about freedom and personal choice enhancing intrinsic motivation of students are also important to note. Lehmann et al. base their argument for this upon evidence which shows that musicians practise pieces of music which they have chosen in a different way to those which the teacher has set. Also, professional musicians took part in more improvisation early in their development than those who did not go on to become professional. The second piece of evidence could be interpreted in a number of different ways. In the book they see improvisation as a way of gaining a sense of choice; however, the willingness and ability to improvise could simply be seen as a sign of higher musical ability and confidence, which could be why the improvisers went on to become professional. On the other hand, intrinsic motivation, by the definition given at the start of this section is about seeing oneself as the reason for one's actions, so it follows that a sense of choice will enhance intrinsic motivation. This could also be the reason why extrinsic motivational factors in music such as parental involvement can have negative as well as positive effects upon motivation in cases where the involvement becomes too controlling.

3.5.3 Motivation – Design Implications

So how can we use the various approaches to motivation outlined when thinking about feedback? Feedback is an external source of motivation and will most likely act to increase extrinsic motivation. There is a danger highlighted by the work of Calder and Staw (1975) that by adding a source of extrinsic motivation, intrinsic motivation could actually be damaged. However, their study looked at a particular case when the task was of little importance to the participants and the reward was not dependent upon them doing it well, only upon them doing it. Other research shows that in the case when the task is considered worthwhile and the reward is seen as signifying achievement the intrinsic motivation can be enhanced (Greene & Lepper, 1974). This is important when designing motivating feedback. Firstly, feedback needs to be seen as a measure of their achievement. This comes from connecting the feedback to how well they are doing a certain aspect of playing. Secondly, feedback must be connected to a worthwhile task. Simply assuming that the students using the feedback consider learning the violin a worthwhile task is not enough. The feedback should also be relevant to their own particular difficulties in learning. For example, feedback on straight bowing is not relevant if the pupil perceives this as easy and finds that holding the violin is their main problem. How feedback is given needs to come from discussion with both pupil and teacher so that the pupil knows that it is relevant and worthwhile for them.

To persist in an activity, students must also feel that success is due to their effort and ability, and that failure is due to a mixture of lack of effort or unstable factors such as bad luck but not due to lack of ability. Feedback needs to support this view of learning the violin, therefore feedback should be appropriate to their level, so that with enough effort the student can succeed in getting 'good' feedback (in some setups this may simply be an absence of 'bad' feedback). Looking at this from an expectancy value point of view the student needs a reasonably high expectancy of success as long as they are willing to put in enough effort. As the student practises it may be the case that they will be able to get ‘good’ feedback with less effort. In this case, it may be that the level of difficulty will need to be increased to maintain the attributional connection between effort and success. In order to encourage an incremental view of musical ability this change in level could be made explicit, i.e. saying to the student in some way ‘by this much practice and effort your ability to do this particular element of playing has gone up a level’.

Lehmann et al.’s (2007) analysis of motivation brings to light aspects which can affect the intrinsic motivation of musical activity. These seem to be rooted in encouraging exploration and a sense of fun, and allowing the pupils to make choices about how and what they learn. Therefore, the design and presentation of real-time feedback needs to aim to support this sense of fun. Moreover, there should be opportunities within the feedback for players to exert some choice over the way they use it. The two models of motivation in music discussed are quite long-term views of what makes an individual choose to persist in playing music. On the other hand, real-time feedback is transitory; it tells the learner to change their playing in the moment. For this reason it is necessary to look at the shorter-term effects of feedback, from one practice to another. The main model that lends itself to this kind of short-term view of motivation is an expectancy x value model that focuses on the motivation to carry out a task. This suggests that feedback may influence motivation either by changing the learner’s expectancy that practice will lead to improvement or by changing the learner’s perception of the value of improving that skill.

3.6 Summary

This review of a range of theories relevant to learning to play the violin and the role of feedback in facilitating this shows that there are many factors to consider. Real-time feedback needs to be able to attract the attention of the player, and for this, exogenous cues will be needed such as flashing or moving visual signals or analogous signals in the tactile modality. It then needs to be perceived and interpreted by the player to enable them to understand movement in terms of the embodied process of playing the violin. Nonetheless, the focus of the feedback should not be too internal as this can limit performance; rather the feedback needs to act as either an external locus of focus or direct attention externally, such as thinking about the body in relation to the instrument and the physical space around it or the musical outcome of the movement.

The choice of modality is also a key design decision. Working Memory, Multiple Resource Theory and Threaded Cognition all suggest that giving feedback in a modality which is not currently used in playing the violin, such as ambient visual feedback or tactile feedback, will minimise the load on players because different sensory modalities can be processed concurrently. However, there is also evidence to suggest that different modalities are suited to taking in different kinds of information and that giving information in an ill-suited modality can be less effective than giving two streams of information in the same modality.

Vibrotactile feedback appears to be a particularly promising type of feedback because, being positioned on the body, it should be well suited to communicating information about physical movement. In addition, reaction time studies show it can quickly catch attention even when other modalities are in use. On the other hand, the movement involved in playing the violin may interfere with perceiving the tactile feedback. This is unlikely as these effects generally only happen when the vibrotactile stimulus is low intensity compared to perceptual thresholds, but it may be that the additional cognitive loads involved in violin playing change this.

Ambient visual feedback is also another possibility for real-time feedback. This appears to have limited bandwidth in terms of communicating information and feedback design needs to use the type of signals that it is best suited to perceiving, such as movement rather than recognising symbols or objects. If more complex information needs to be communicated through visual feedback, then ambient feedback could be used to catch attention, and focal vision could take in the detail. However, this will draw the eye away from the music. In either case, it will be best if visual feedback is positioned close to the music as ambient vision is best close to the line of focal vision.

Ambient visual feedback does have some possible advantages over vibrotactile feedback because it may be better suited to giving motivating feedback. For young players to persist in playing the violin, lessons and practices need to be fun and intrinsically motiving. Vibrotactile feedback is limited because it only has a restricted palette of sensations, for example rough or smooth vibrations (Hoggan, Brewster, & Johnston, 2008). The body centred reference frame may make it good at encoding movement information but is not necessarily very expressive in other ways. For example, a theatre company who developed an immersive performance using haptics (van der Linden et al., 2011) chose not to use vibrotactile feedback because they felt it was “strangely empty” and “cold”. On the other hand, visual feedback can be more expressive because it can come in many colours and patterns. This gives it more scope for being designed to be motivating and enjoyable to use.

The research in this thesis is based on previously published research that investigated real-time vibrotactile feedback for learning the violin (van der Linden, Schoonderwaldt, Bird, & Johnson, 2011) (van der Linden, Johnson, Bird, Rogers, & Schoonderwaldt, 2011). This is outlined before proceeding to the methodology chapter, as the findings from these studies informed the in-the-wild approach to the research adopted in this thesis.

3.7 MusicJacket: Initial Research on Using Real-time Feedback and Violin Playing

At the beginning of my PhD, I collaborated with a group of researchers at the Open University to conduct research that focused on how real-time vibrotactile feedback could be applied to learning the violin. In consultation with violin teachers we designed and built the MusicJacket - a suit that uses an inertial motion capture system to measure the bowing trajectory and violin hold of a player and gives vibrations on the arms and torso to indicate when they are making a mistake. This was tested and improved upon iteratively until it took the form of the design shown in Figure 3-4. The full design process that went into the MusicJacket is given in Appendix A. This was based on a motion caption system intended for the entertainment industry but was appropriated here to monitor arm and body movement and to provide real-time feedback based on deviations from an ideal movement. Multiple Resource Theory (Wickens C. , 2008) led us to choose vibrotactile feedback as this modality appears to be used less than vision and audition when playing violin. Ideas from Embodied Interaction (Dourish, 2001) and an aim to reduce the Gulf of Evaluation (Norman, 1986) also influenced our choice as we believed that giving the feedback on the body would enable players to understand how the feedback related to their movement more easily. Vibrators were initially positioned according to the push metaphor described by Spelmezan et al. (2009). However, as the design progressed we improved upon the position of the vibrotactile feedback in consultation with violin teachers and an Alexander Technique expert ((Artist in Balance, 2012), (Society of Teachers of the Alexander Technique, 2011)) and based on the results of a small scale user study (Johnson, van der Linden, & Rogers, 2010). In this process we found that vibrators positioned on the muscles in the arm made the muscles feel like they should contract, and were also reported to be uncomfortable. We also found positioning many vibrators on the same bone in the arm made it hard for participants to differentiate between them because bone conducts the vibration. This problem was solved by positioning the vibrators near joints, which are also the anatomical reference points described by Cholewiak and Collins (2003) in their studies of tactile perception on the arms. The Alexander Technique specialist also advocated including vibrators on the body to encourage good upper body posture as well as arm position. Different positions for vibrators were tested and it was found that vibrators positioned lower down on the torso were uncomfortable as they caused the diaphragm to vibrate, whereas positioning them on the ribs gave a sensation of lifting up.

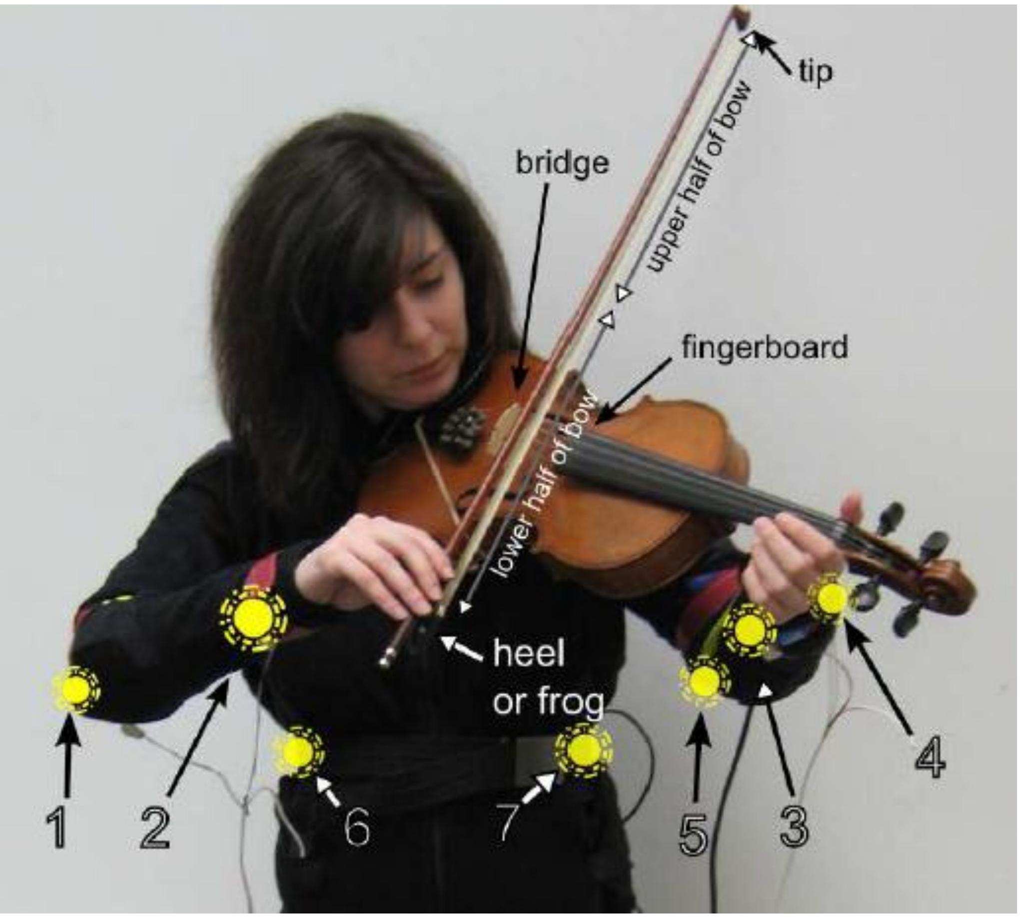

Figure 3-4: The MusicJacket (1 move bow hand forward, 2 move bow hand back, 3 move violin left, 4 move violin right, 5, 6, 7 lift violin)

Two studies were conducted with the MusicJacket to investigate the efficacy of real-time feedback for learning the violin. The first study was a laboratory study with eight complete beginners. Four participants used the MusicJacket and four practised without for two out of six session spread over a week. The motion capture jacket was used to measure differences in the participants' playing. Using this method we were able to show that the MusicJacket did significantly improve the way people bowed compared to their normal bowing behaviour. However, we were not able to show significant differences between the control and test groups by the end of the week when the participants' played without the feedback. We also noticed extreme differences between participants over the week even though they were all complete beginners at the start of the study. Some participants found straight bowing a very difficult task to learn in the initial lessons and the feedback was able to help them in particular, whereas others did not find it such as challenge. Body shape and approach to music and physical learning are all possible causes for these individual differences. This made it difficult to compare the test and control groups because each participant had different learning needs. A report of these findings has been published here (van der Linden, Schoonderwaldt, Bird, & Johnson, 2011) and a more detailed report of this study can be found in Appendix B.

The second study of the MusicJacket took place in violin lessons with violin teachers and their pupils. This study ran over five sessions spread over two months. There were two teachers; one used the MusicJacket with children in their first 6 months of playing aged between 6 and 9, and the other used the MusicJacket with children who had been playing two or more years aged between 10 and 15. The focus of this study was how real-time vibrotactile feedback changed normal violin lessons and whether this led to improved learning. We found that the way real-time feedback impacted on learning was specific to the child's learning needs. The younger children had difficulty holding the violin and in these cases we were able to improve their violin hold using the MusicJacket (see Figure 3-7 and Figure 3-8 for examples of good and bad posture). Some of the older children had difficulty with straight bowing and in these cases, we were able to show that the Musicjacket helped them to improve their bowing by making them more aware of it. One of the older children had no problem with either of these aspects of technique but had difficulty using the full length of the bow when she played. Using a full bow means playing from close to the heel of the bow to near to the tip (see Figure 3-5 and Figure 3-6 for terminology). In the final session, the Musicjacket was redesigned to support her in her learning goals by using vibration as positive feedback when she played using enough bow. This resulted in a large improvement in her bow use over the session where she used this new specialised version of the MusicJacket.

The in-the-wild nature of this study was key to the findings it produced. For example, when participants were playing particularly demanding pieces of music they reported times when they felt no feedback, but the logs show that vibrotactile feedback was given to them during that time. This was not found in the laboratory study. It was also observed that the older children and their teacher used the vibrotactile feedback as a way of discussing movement, improving communication between student and teacher. This study initially used a between participants study design with a test group which used the feedback during lessons and a control group which had lessons without feedback. However, as the study progressed we chose to change the study design so that all participants had feedback. This was because the students were all very different from one another in their abilities, personalities and learning goals, and the teachers adapted the way they taught the violin to the individual needs of the students. Therefore, we felt it would not be valid to directly compare the quantitative results from the two groups. Instead, we took each participant as a case study of how real-time feedback can contribute to teaching individual learners. The findings from this study are published here (van der Linden, Johnson, Bird, Rogers, & Schoonderwaldt, 2011) and a full report of this can be found in Appendix C.

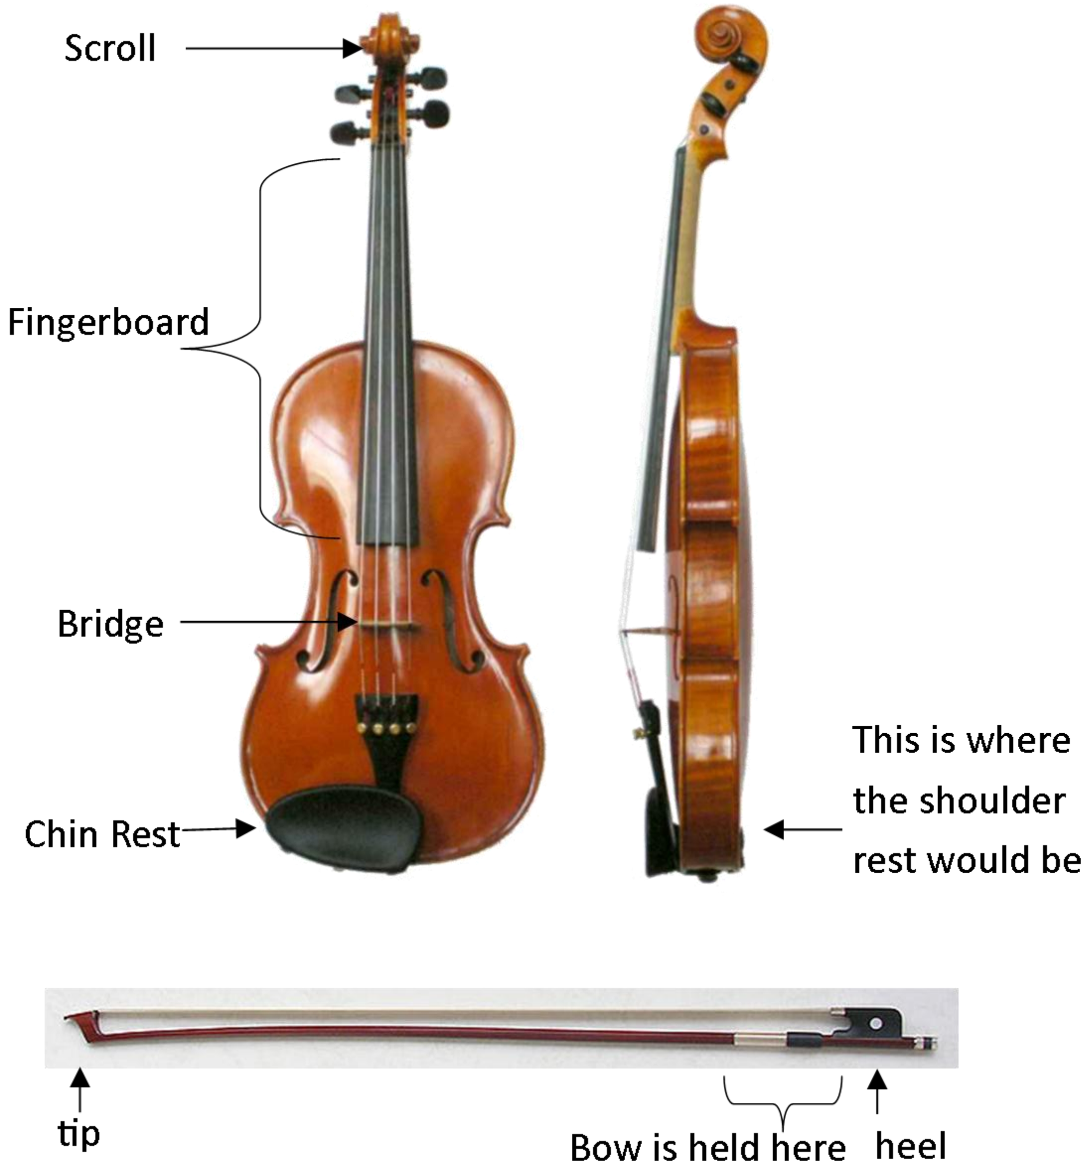

Figure 3-5: Parts of the violin. Picture of the violin taken by 'just plain Bill' and used under creative commons license CC0 1.0. Picture of the bow taken by Georg Feitscher and used under creative commons licence Attribution-ShareAlike 3.0

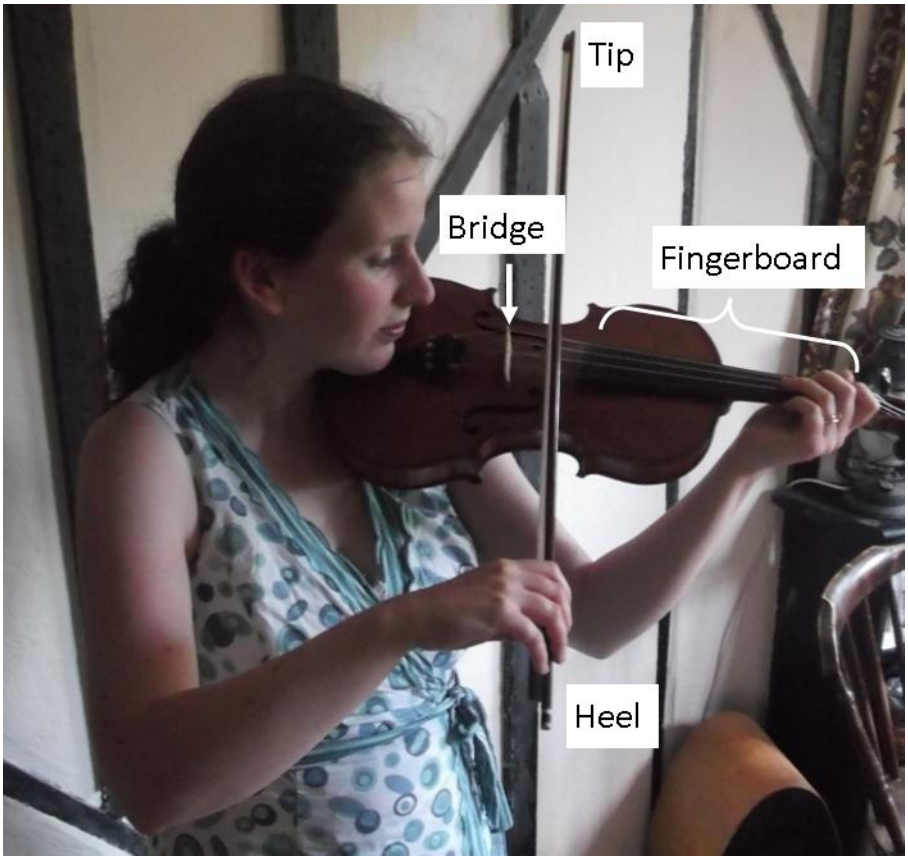

Figure 3-6: Parts of the violin in use.

Figure 3-7: Correct hold of the violin with scroll pointing out horizontally

Figure 3-8: Incorrect violin hold, with scroll pointing downwards

3.7.1 MusicJacket Conclusions

From a design perspective the development of the MusicJacket suggested some key factors to consider when people play the violin and how these could be explored when trying to improve playing in a variety of contexts. These form the basis of a series of studies that are reported in this thesis. First, the two learning foci for prototypes in this thesis draw on aspects of violin technique that were discovered to be important when studying the MusicJacket. These are bow length and violin hold/posture. Using a full bow means playing from close to the heel of the bow to near to the tip. This is a desirable way to play as it makes a richer sound. Not all pieces require full bows but it is important to be able to play them when necessary and for some learners this can be very difficult. Holding the violin up and keeping good posture is also important and can be difficult for learners. Good posture and violin hold helps the violin to resonate and makes it easier to play in tune and bow correctly.

The development of the MusicJacket also informed where best to position the vibrators to be effective. One implication is that it is important to position vibrators with respect to the underlying anatomy of the player. Positioning vibrators on bone rather than muscle was found to be more easily interpreted and more comfortable. Positioning vibrators close to anatomical reference points such as joints, and minimising the number of vibrators per bone, aids perception and interpretation. Studies with the MusicJacket in the design stage also showed that having feedback about bowing and posture at the same time was overloading and confusing for participants. Therefore, a single learning focus for real-time vibrotactile feedback was adopted for the prototypes in this thesis.

The studies with the MusicJacket were also informative from a methodological point of view. Both the lab study and the in-the-wild study showed strong individual differences in the way participants played and learned. These were difficult to take into account when comparing control and test groups and we found it more informative to focus on how the participants' playing changed between different conditions rather than analysing the differences between groups of participants. The in-the-wild study also brought out many more interesting findings which were not present in the lab study. These come from the study being held in a more naturalistic environment where participants are carrying out realistically demanding tasks while also having to attend to their teacher's instruction and relate them to their own personal learning goals. This is a more challenging environment to carry out research in but also a more interesting and ecologically valid one for the study of real-time feedback for learning to play the violin. All the studies reported in this thesis take an in-the-wild approach where the goal is to study real-time feedback with participants in settings that are ecologically appropriate.

Key Points

- Posture and long bowing are useful learning goals which some learners find difficult
- It is important to position vibrators with respect to the anatomy beneath the skin
- Only give feedback about one learning focus at a time to avoid confusion

- There are strong individual differences between learners in the way they play and learn the violin. This makes it difficult to draw a fair comparison between players in a between participants study design with a small sample size.
- In-the-wild studies of violin playing produce different findings to lab based studies. Findings from in-the-wild studies are more likely to be relevant to designing real-time feedback for learning violin, as learning violin is a situated activity.

Chapter 4 – Methodology

In light of the findings from the initial MusicJacket studies, I chose an approach that would support studying violin playing in-the-wild. Learning to play the violin may take many different forms, for example, individual practice, taught lessons or learning to play in an ensemble. The role real-time feedback can usefully play may be different in each of these learning scenarios. This means that each setting should be investigated separately as there will be a different set of research questions and design requirements for each type of learning.

There are also a number of approaches to measuring and understanding how real-time feedback affects learning. One metric for measuring learning is using sensor technology to quantitatively measure physical changes in a learner's playing. However, facilitating learning can be conceptualized in other ways as well. For example, improved communication between teacher and pupil is a way in which real-time feedback may aid learning. Motivating learners is another example of a way in which real-time feedback can help learning. There are other ways that the effects of real-time feedback on learning to playing can be qualitatively measured. This thesis uses both quantitative and qualitative research methods to investigate real-time feedback. Moreover, the interpretive framework used to understand how real-time feedback influences learning evolves from an experimental framework to a qualitative ethnographically informed approach as the research questions change from focusing on learning as a physical change in playing to a more holistic view of learning as a long term process.

A goal of this thesis is to understand how different kinds of real-time feedback can affect different aspects of learning by applying it to the practical example of learning the violin *in-situ*. Within this perspective, design and research are viewed as intimately linked. The methodology described here addresses the importance of the design process as well as addressing the research questions posed in the thesis. As the perspective on learning evolved through the programme of research so did the interpretive framework for understanding the findings. In so doing, the design and research questions evolved simultaneously.

The methodology is structured in four layers. These describe the multiple ways in which the research approach was adapted through the process of design (Table 4-1). The layers are:

(i) Evolving research questions and design iterations

(ii) Setting

(iii) Interpretive framework

(iv) Data collection methods.

Each of these layers is described in more detail below. The way they changed through the course of the research and were used to inform each study is then outlined. A more detailed description of the methodology for each study is given in the subsequent chapters, themselves (chapters 5-7).

Layer 1 – Iterative Design and Evolving Research Questions	**Research Questions**	**How do visual and tactile modalities of real-time feedback differ in supporting individual violin practice?**	**How can both shared and personal real-time feedback support playing as part of an ensemble?**	**How can real-time feedback fit into the extended process of learning through various forms of practice?**
	Design Iteration	1. MuSense Real-time vibro or visual feedback about bowing	2. ShareSense Real-time vibro & shared visual feedback about bowing	3. Buzzy Jacket and Twinkly Lights Real-time vibro and visual feedback about posture
Layer 2 – Setting	In-the-wild Approach	University Practice Rooms	Musical Summer school	At School and at Participants' Homes
Layer 3 – Interpretive Framework	Experimentally Informed			
	Ethnographically Informed			
Layer 4 – Data Collection Methods	Motion Capture Sensors			
	Observations			
	Interviews and Comments			
	Video Analysis of Sessions			
	Audio Recordings of Sessions			
	Questionnaires			

Figure 4-1: How the methodology was adapted to address the evolving research questions. The intensity of the colour represents the how much that method was used in each study. The hue of the colour is not significant.

The first layer is the overall research structure that guides the inquiry into real-time feedback. This is an iterative approach in which new research questions emerge from the findings of the user studies and are pursued in the next iteration of the design. The research began by focusing on the efficacy of different types of real-time feedback for improving a specific aspect of technique (Q1). Gradually, over the course of the research, the questions evolved to view violin learning more holistically culminating in a study which investigated all forms of practice over a longer period of time (Q3). Concurrent with this, is an iterative design process where the design of the prototypes and user study is guided both by the research questions and the findings from the previous studies. Hence, the two design processes evolved in a mutually dependent way to determine what type of feedback is appropriate in which types of situation.

The second layer refers to the setting used to investigate real-time feedback to aid learning the violin. Many of the studies in the literature which investigate real-time feedback (e.g. (Lieberman & Breazeal, 2007), (Bloomfield & Badler, 2008)) use a lab study to show the efficacy of their system. However, due to their controlled nature we would argue that lab studies are quite different from the real-world and so the findings may not be relevant to how people play and are motivated in their day-to-day lives. People do not live in labs; they live in a rich environment, full of distractions and hidden motivations. People normally make their own decisions in a situated way based on different factors at the time, whereas in a lab everything is controlled and participants expect and are expected to behave according to the instructions of the researcher. Therefore, in-the-wild studies were conducted to understand how real-time feedback can be used to aid learning in the real-world.

In-the-wild studies aim to deploy a prototype in a real-world situation where the prototype is intended to be used. Some of the in-the-wild studies designed take an experimentally informed, hypothesis testing approach; whereas others take a more explorative, qualitative approach. The aim of the latter was to overcome preconceptions about what might be found and observe what happens as people and technology interact.

4.1 Interpretive Framework: Experimentally Informed versus Ethnographically Informed

The third layer is the interpretive framework which is concerned with how to understand the data that are collected. In general, in-the-wild studies are not linked to a particular interpretive framework. In the literature, a variety of methods are used to analyse and interpret data in-the-wild: from statistical tests on quantitative sensor data to test a predicted hypothesis (Rogers, Hazlewood, Marshall, Dalton, & Hertrich, 2010); to qualitative analysis which aims to describe how people interact with technology in that context (Marshall, Morris, Rogers, Kreitmayer, & Davies, 2011).

The two frameworks used in this thesis are broadly termed as experimentally informed and ethnographically informed. Experimental frameworks are used in chapter 5 in order to answer questions concerning the efficacy of particular types of real-time feedback for improving measurable aspects of playing technique. An ethnographically informed framework is used in chapters 6 and 7 to understand qualitatively how real-time feedback is used by different groups of players in a situated way.

An experimental framework interprets data by testing hypotheses and using statistics to show the presence or absence of a link between two or more variables. In the first study in chapter 5 we aimed to show quantitatively that the presence of different types of real-time feedback causes a measurable improvement in playing. Collecting and interpreting data in an experimentally informed way has the advantage of producing findings which combine data from many participants to give a strong evidence base and a small probability of false findings. But important findings may be missed because they were not part of the hypothesis being investigated. For example, we may hypothesise that real-time feedback will cause learners to improve the way they play one set of exercises which we have chosen to measure, but it maybe that the more interesting data is in the way the participants apply the real-time feedback to a different set of pieces with different challenges. A more explorative method might be able to reveal these more.

If the outcomes of a study are unpredictable then it is difficult to formulate hypotheses before the study begins that will still be relevant and meaningful to the study after data have been collected. Since interpreting data in an experimentally informed interpretive framework is dependent upon hypotheses this can make it very difficult to apply it to interpreting the data from these unpredictable studies.

An experimentally informed interpretive framework also tends to treat participants as the same, but there may be individual differences that, from the participant's point of view, can affect their experience of real-time feedback. This is often overlooked in the way data is statistically analysed when testing a hypothesis and drawing conclusions. Moreover, an experimental approach tends to prioritise studying factors that are quantitatively measurable; however, these may not be the most important factors influencing the study.

An experimentally informed interpretive framework can be useful when conducting studies in-the-wild. However, it does not capture everything that is happening in-the-wild, particularly factors arising from the social context and individual differences in the way participants approach learning. Therefore, in addition, a more qualitative approach is needed to investigate and understand the contextual factors. Ethnography is a method used in HCI to understand people, their activities and the context of a particular setting. It is not usually used for prototypes deployment. However, when studying technology in-the-wild, the interpretive framing from an ethnographic perspective has the potential to provide new insights.

In ethnography, researchers typically immerse themselves in the setting over a period of time, collecting data through observation, participation and informal interviews. The aim in this programme of research was to learn to understand the world from the perspective of the context being studied. Often participant observation is used to enable the researcher to transform their point-of-view. It gives researchers an insider perspective, allowing them to interpret their observations with an intimate understanding of the values, beliefs, social norms and experiences of the group or culture being studied. Another key element of the ethnographic approach is reflexivity - reflecting upon the way the ethnographer is part of the ethnography itself, their background and perspective is as much part of the research as those of the culture being studied because it governs the way the researcher interacts with participants and how they interpret participant responses ((Rode, 2011), (Dourish & Bell, 2011)).

In the ethnographically informed interpretive framework used in the later studies of this thesis, participant observation and reflexivity were used to find an authentic perspective from which to interpret qualitative data such as observations, video analysis and interviews and to draw conclusions about how the social context and individual differences influence the use of real-time feedback.

4.2 Data Collection Methods

The fourth layer of the framework refers to the specific set of methods chosen to collect data from a particular study. Many methods can be used for in-the-wild studies to collect data, such as: sensor readings (Rogers, Hazlewood, Marshall, Dalton, & Hertrich, 2010), observations (Marshall, Morris, Rogers, Kreitmayer, & Davies, 2011), video analysis (Hinrichs & Carpendale, 2011), interviews (Rogers, Hazlewood, Marshall, Dalton, & Hertrich, 2010), focus groups (Jambon & Meillon, 2009), surveys (Rogers, Hazlewood, Marshall, Dalton, & Hertrich, 2010) and diaries (Johnson, van der Linden, & Rogers, 2010). This shows how in-the-wild user studies are not attached to a particular set of methods, but utilise many of the methods already used in HCI research to analyse the interactions between users, technology and setting with an emphasis on capturing how the setting influences interaction.

A key feature of all the in-the-wild studies referred to here is that many of the findings were unexpected, and not predicted by the researchers before the study commenced. The findings from the lab studies do not necessarily apply in-the-wild, making it hard to use these to formulate the predictions or hypotheses which would normally be used to guide study design. This demonstrates both a challenge and an advantage for in-the-wild studies. The advantage is that in-the-wild studies provide a different perspective on technology to lab studies and may reveal previously undiscovered phenomena and insights. The challenge is that it is important to choose a set of methods that will be able to capture unexpected findings; qualitative methods such as observing and videoing may be best suited to this because they are the most flexible for recording unforeseen happenings. This challenges us to adopt a creative approach; often involving using mixed methods to collect data in order to reveal unforeseen findings.

A number of different methods were used for collecting data, including sensor data to interviews and observations. Using several methods in one study can confirm findings through triangulation or to provide complementary findings to give fuller picture of how real-time feedback is used and experienced. Table 4-1 briefly summarises the data collection methods used in this thesis and the main purposes they were used for. Figure 4-1 shows which methods were used in which study. In the next section we describe the methods used in each individual study in more detail with the aim of giving an overview of how the methodology evolved in response to emerging research questions. Every chapter also has a methodology section which justified the choice of methods for each particular study.

Method	**Use**
Motion capture	Quantitatively comparing playing movement/posture in different experimental conditions. Visually representing movement so that it can be discussed qualitatively
Video	Gathering information about whole body movement (rather than focusing parts of the body being sensed) Keeping a detailed record of events that occur in study sessions
Observation	Gathering information about whole body movement (rather than focusing on parts of the body being sensed) Gathering information about social interactions and interactions with the setting Gathering unexpected findings
Interviews	Gathering participant opinions about the prototypes Gaining insight into participant motivations and goals Gathering descriptions about the participant experiences using the prototypes
Participation	Gaining a first person perspective on the experience of using real-time feedback in the study setting Building rapport with participants and gaining an insider perspective on being a member of the ensemble
Questionnaires	Gathering participant opinions on the prototypes quickly and uniformly

Table 4-1: Summary of research methods used in this thesis

4.3 Methods Used in Each Study

As the research questions evolved from one study to the next, the methodology was adapted and reformulated in order to answer them appropriately. In addition, with each user study, lessons were learnt about the methodology itself. For example, from the first study emerged the importance of individual differences in the way people learn and use feedback. This led us to change our interpretive framework from an experimentally informed one which treats the data from participants identically, to an ethnographically informed one. Using an ethnographically informed framework enabled data from each participant to be interpreted in terms of the researcher's knowledge of that individual built through joint participation and the building of rapport. These methodological lessons are discussed in more detail at the end of each study chapter (chapters 5-7) and also in the methodological reflections in chapter 10.

Study 1: How do visual and tactile modalities of real-time feedback differ in supporting individual violin practice?

The aim of the study described in chapter 5 was to compare two modalities of feedback: vibrotactile feedback and visual feedback. An in-the-wild study was carried out where university orchestra players practised their current orchestra music in university practice rooms. The task and setting were familiar to the participants. Being on the university site also gave some control over the study setting and having participants that were all working on the same music meant that it was possible to take an experimental approach. A within participants method was used where participants practised with different types of feedback and their playing was compared statistically. Sensors were used to measure how different conditions affect their playing. As well as the experimental interpretation of results a more qualitative case by case approach was adopted. Interviews were used to give insight into what it is like for participants in the different conditions. This approach was also informed by the researcher's pre-existing relationships with the participants as a member of their orchestra.

Study 2: How can both shared and personal real-time feedback support playing as part of an ensemble?

The aim of the study presented in chapter 6 was to investigate the experience of using real-time feedback while playing in an ensemble. An in-the-wild study was conducted with ensemble players at a summer school. An ethnographically informed approach was taken where I actively participated as a violinist in the study and the study setting. By participating in this study I was able to learn about the experience of being an ensemble player and build rapport with the study participants. Videos were also used to record sessions. A questionnaire after the end of the summer school was employed to obtain opinions about the system. The participative experience enabled me to interpret the observations and responses from the participants.

Study 3: How can real-time feedback systems fit into the extended process of learning to play the violin through various forms of practice?

The aim of the study described in chapter 7 was to investigate how children use feedback themselves in their day-to day practice settings over many weeks, including ensemble practice and home practice. This study was carried out in-the-wild. During the ensemble practice session I was present in the role of a teacher, running the sessions while making observations. The children were asked to keep a video diary. This was combined with an interview with each child at the end of the study to find out about using the system at home and at school.

4.4 Summary

The methodology employed in this thesis was to conduct a series of in-the-wild studies. For each study, the aim was to choose the most appropriate methods to answer the specific research question addressed. Methods were chosen based on their suitability for the questions and setting and their efficacy in previous chapters. This iteration leads to trying out new methods for prototype evaluation, such as participant observation. It also enabled me to continuously reflect on the methodological approach adopted and underlying assumptions. Throughout this process of reflection, new understandings about the methodology emerged.

The next chapter describes the design of and study of MuSense - a prototype device which could give either vibrotactile or visual real-time feedback. The purpose of this prototype was to compare the visual and vibrotactile sensory modalities as ways of giving feedback by investigating which modality was most effective for improving bowing.

Chapter 5 – MuSense: Comparing Visual and Vibrotactile Feedback

5.1 Introduction

Our initial studies of the MusicJacket (briefly reported in section 3.7) (van der Linden, Schoonderwaldt, Bird, & Johnson, 2011), (van der Linden, Johnson, Bird, Rogers, & Schoonderwaldt, 2011) investigated how effective vibrotactile feedback was in facilitating beginners learning to play the violin. Vibrotactile feedback was chosen as the modality in which to provide feedback so that it did not interfere with the main visual and auditory processes that playing the violin involves. Force feedback was not investigated because the research literature suggests that passive movement is not as effective for motor learning as voluntary actions, as discussed in the literature review. Overall, the findings from the research with the MusicJacket showed that this form of real-time vibrotactile feedback was effective at helping beginners learn to play the violin. In particular, the findings showed that it improved long bow playing by vibrating to show when the student reached the end of their bow to indicate that they had used the full length of the bow. For one of the students who found using full bows difficult this feedback transformed her playing (van der Linden, Johnson, Bird, Rogers, & Schoonderwaldt, 2011).

However, the studies also showed that vibrotactile feedback can sometimes interfere when playing the violin and, conversely, concentrating on playing the violin can interfere with the extent to which the vibrotactile feedback is perceived by the person. When students used the MusicJacket, they often made more mistakes in the piece of music they were playing; and when students were playing particularly difficult pieces they sometimes reported not feeling any vibrotactile feedback even though it was active (van der Linden, Johnson, Bird, Rogers, & Schoonderwaldt, 2011). This tension of being able to perceive and act upon vibrotactile feedback when engrossed in an intense sensory-motor activity raises the question of whether visual forms of feedback might actually be more effective – provided they are distinctive, and are able to attract the person's attention but do not interfere with the on-going motor task. In particular, the visual modality could provide real-time information to inform the players of their progress in reaching a specific goal – which is more difficult to do with vibrotactile feedback. However, it needs to be distinct enough from the music score that is being read when playing so that it does not interfere with it.

Visual feedback has a number of features that contrast with vibrotactile feedback. For example, the player has a larger degree of choice over whether they use the feedback – they can choose to look at it, or choose to ignore it. In contrast, they cannot consciously control whether they sense the vibrotactile feedback. Visual feedback also has potential for depicting greater complexity of information since humans can see with higher resolution than they can feel tactile vibration. Theories of learning and motivation (Dweck & Leggett, 1988) suggest that providing information that allows someone to see their incremental progress more clearly will motivate them to work towards their goal. Additionally, information can be encoded in colour as well as space, which vibrotactile cannot use - although the latter can be varied in intensity and pulsing in a similar way to visual representations. On the other hand, vibrotactile feedback is presented to players on their body, using the body as a frame of reference; this is not possible with visual feedback. This raises the question of how 'being touched' rather than 'looking at' affects the way people understand and interpret feedback.

The first study presented in this thesis addresses the first research question from section 1.3 which concerns the efficacy of providing different forms of real-time feedback whilst practising to play the violin:

How do visual and tactile modalities of real-time feedback differ in supporting individual violin practice?

To answer these questions, a study was designed that compared vibrotactile feedback with visual feedback for learning to practice long bowing. Long bowing is an aspect of violin playing technique where players use close to the full extent of the bow length for each note. This gives a richer sound than using only a short length of bow. This action is an important part of playing the violin well that learners often find difficult. If they get into the habit of only making small movements with their bow arm, then it is very difficult for students to overcome this without feedback as we observed when studying the MusicJacket (van der Linden, Johnson, Bird, Rogers, & Schoonderwaldt, 2011). As players progress they need to be able to match the length of bow that they use to the style of piece they are playing, this includes long bows for smooth flowing music as well as short bows for fast or staccato (spikey) music.

5.2 Hypotheses

Two hypotheses were formulated based upon the discussion of the theories ((Wickens, 1980, 2008), (Dourish, 2001), (Dweck & Leggett, 1988)) presented in the literature on the efficacy of using different modalities for real-time feedback:

(i) Vibrotactile feedback will be more effective at changing long bowing compared with visual feedback.

(ii) Visual feedback that shows progress towards a goal will be more effective at changing long bowing compared with visual feedback that only shows when a goal has been reached.

The first hypothesis is based on two lines of argument. Firstly, previous research has shown that vibrotactile feedback is more effective than visual feedback for influencing movement (Bloomfield & Badler, 2008). Secondly, from a theoretical perspective, vibrotactile feedback is predicted to be more effective than visual feedback because the visual channel is already heavily loaded with the task of reading music whereas the tactile channel is not (Wickens, 2008). Moreover, vibrotactile feedback uses a body-centred reference frame which should make it easier to interpret than the visual feedback which is more abstract.

The second hypothesis investigates whether different forms of visual feedback are more effective at showing the player that they are approaching their goal. It was hypothesized that visual feedback that shows the build up towards a goal will be more motivating than one that provides feedback only when a goal has been achieved (Dweck & Leggett, 1988). A more complex form of vibrotactile feedback was not included in the study because studies with the MusicJacket (Appendix A) had suggested that it was difficult to notice multiple vibrators while bowing the violin. Hence, the study compared three conditions: (i) vibrotactile and (ii) simple visual and (iii) complex visual. The first visual condition (labelled simple) flashes when a long bow has been reached, and the other condition (labelled complex) was designed to show the build up towards the goal of reaching a long bow. Having more continuous feedback from the in the complex condition should make a clearer connection for the player between movement and feedback.

Based on Benford et al.'s (2005) framework for designing sensor systems, the feedback focuses on the *"desired"* movements that correspond with ones which can be *"sensed"*. These movements will be *"expected"* by the user because they are part of the normal action of playing the violin. However, whether long bowing is an *"expected"* movement to get feedback about will depend on whether the player values long bowing as an aim for improving his or her playing.

5.3 Design

The prototypes developed for this study were intended to be tested in real-use practice settings. Specifically, the focus was on how feedback could be used to improve independent practice. Building devices that can be given to participants to be used independently, without researchers present, poses a number of demanding design criteria, including the following:

- **Robustness:** The prototype must be able to withstand day-to-day use without getting broken.
- **Price:** There is more chance that prototypes can be lost or broken; therefore, they cannot be too costly.
- **Set-up:** The prototypes must fit into daily activities so they must not be time consuming to set-up and use. Ideally, they should just be able to be switched on and off.

- **Ease of use:** Prototypes must be easy to use and understand.
- **Reason for use:** The participant must have a genuine motivation to use the prototype in their daily practice. They need to be able to understand its relevance to their learning and be able to see improvement whilst using it.
- **Single focus:** Feedback should focus on one aspect of technique at a time otherwise students get confused.
- **Personalised** feedback: Feedback should relate to the learning goals presented to the student by the teacher.
- **Flexibility**: A prototype should be flexible enough to change the focus of the feedback through the course of the study if the student's goals change.

Further criteria are based on the research questions comparing the visual modality with the tactile modality.

- **Comparing feedback modalities:** It should be possible to give feedback about the same aspect of technique using different modalities (i.e. there should be a visual setting and a vibrotactile setting which must in other ways be equivalent)
- **Recording data:** The device may also need to offer some way of recording sensor reading to enable analysis of whether the feedback is affecting the way participants play.

Additional criteria come from the physicality of playing the violin.

- **Freedom of movement:** The device should allow the arms and body to move freely, meaning devices should be lightweight and wireless.

Based on these criteria, the initial prototype was designed to be affordable, robust and use a minimal set of sensors. It was also designed to be used without a laptop computer. Instead all the sensor data was processed on a small microcontroller. This helps to keep the system low cost and makes it more robust and quicker to set up. However, the single sensor set-up can potentially limit the ability to *personalise* the system or use it *flexibly* to move from one learning task to another. To address this, the final system was designed to be modular with several sensing devices that each address a specific aspect of technique, and several feedback devices which each give feedback in a different way (Figure 5-1). Each sensing device can be paired with a feedback device so that any style of feedback can be given about any element of technique. To enable the prototype to be untethered, the modules communicate using radios and are powered by batteries. The issue of batteries poses one particular practical problem as a balance needs to be made between the capacity of the battery and its weight.

The prototype used in the user study consisted of a sensor band to measure bow usage, a vibrotactile band to give vibrotactile feedback, and an LED display which could show one of two types of visual feedback above the music. How each one works is explained below.

MuSense

SENSOR DEVICES

SENSOR: Bow Length

This sweat-band that is worn on the lower arm away from the wrist. It is designed to infer the length of bow the player uses by measuring the angle which their arm moves through as they play.

It does this by integrating the reading from a gyroscope. The bow length sensor can be wirelessly coupled with two feedback devices—the LED line and the vibrotactile band.

THIS SYSTEM COULD BE EXTENDED WITH MORE SENSORS. For example...

LEFT-ARM POSTURE SENSOR

This device could sense whether students have unconsciously dropped their violin down whilst playing.

MUSCLE TENSION SENSORS

A device could be worn to sense whether players were inappropriately tensing their muscles.

PLAYERS CAN CHOSE THE SENSOR DEVICE MOST RELEVANT TO THEIR PRACTICE

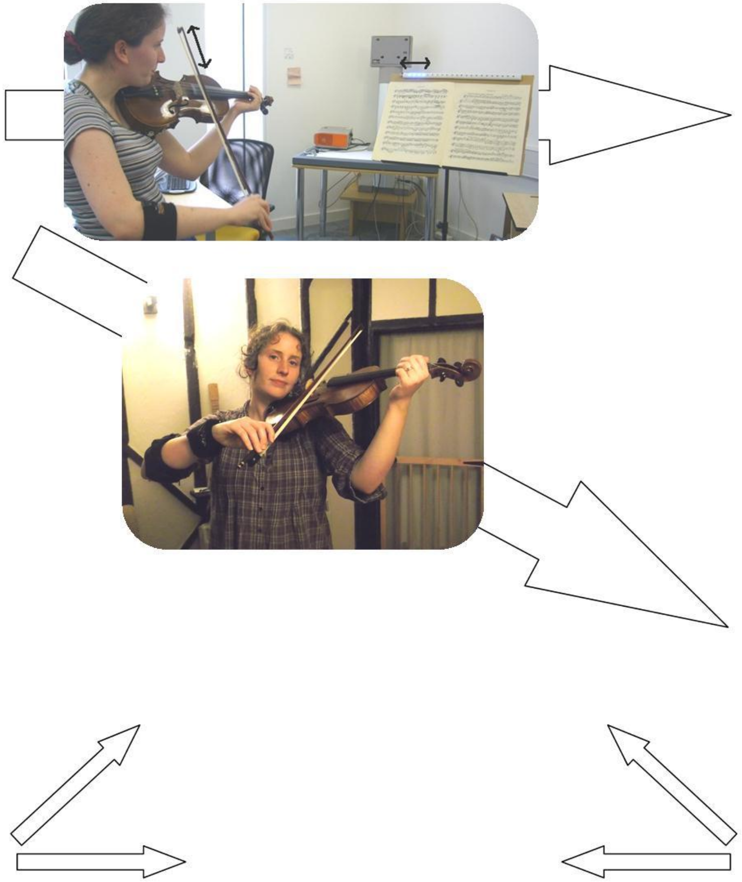

FEEDBACK DEVICES

FEEDBACK: LED line

This display is formed from 22 colour changing LEDs mounted in a line above the music. It has two settings, complex and simple. The complex setting follows the movements of the player—the lights on the display build up to match the length of the student's bow stroke. The simple setting just gives an on/off response to indicate whether the student is doing well or not—it lights up green once the play has used a 'good' amount of bow.

FEEDBACK: Vibrotactile

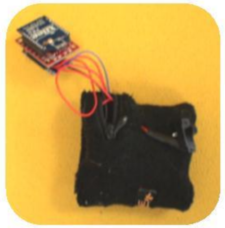

A vibrator is built in to a sweat band positioned on the right wrist and buzzes inform the player when they have used a 'good' amount of bow (full length in this study).

THIS SYSTEM COULD BE EXTENDED WITH MORE FEEDBACK. For example...

ALTERNATIVE VISUAL DISPLAYS

ALTERNATIVE POSITIONS FOR VIBRATORS

ALTERNATIVE MODALITIES

PLAYERS CAN CHOOSE THE FEEDBACK THEY LIKE BEST

Figure 5-1: MuSense System Overview. The Bow Length sensor, LED line and Vibrotactile feedback were built and used in the study in this chapter.

Sensor Band

To measure bow usage the sensor band uses a gyroscope (LISY300AL) which measures rate of change of angle around a single axis. This was sewn into a sweat band so that it could be positioned on the top of the forearm (Figure 5-1, Figure 5-3). The measurements from the gyroscope were read by an Arduino Pro Mini and integrated in order to get the angle which the arm is moving through as the player bows. The measurement was then wirelessly transmitted to the feedback modules using an XBee series 1 radio which was connected to the Arduino Pro Mini.

When using a gyroscope sensor there is a danger that errors build up over time as the signal is integrated. To avoid this happening, each time the direction of the arm changed (i.e. a new bow was started) the gyroscope would begin integrating from zero again because the sensor band only needs to calculate the length of each bow not the angle of the arm over longer periods of time. This made the angle measurement consistent and robust over time.

This sensor band only measures the change in angle of the arm rather than directly measuring the length of bow being used. These two quantities are very closely related (Figure 5-2) but not identical. For example, the sensor could be fooled by twisting from the torso, thus changing the angle of the arm without moving the bow. However, given our approach, which was to make it light-weight and which could be easily used in-the-wild, this level of accuracy was considered adequate for its purpose. Particularly because participants were musicians who knew about bowing and were motivated to use it for improving their bowing rather than simply getting good feedback. Moreover, players could see a clear link between their bow use and the feedback they received. The sensor band was tested with the cello before the study and found to work equally well for this instrument.

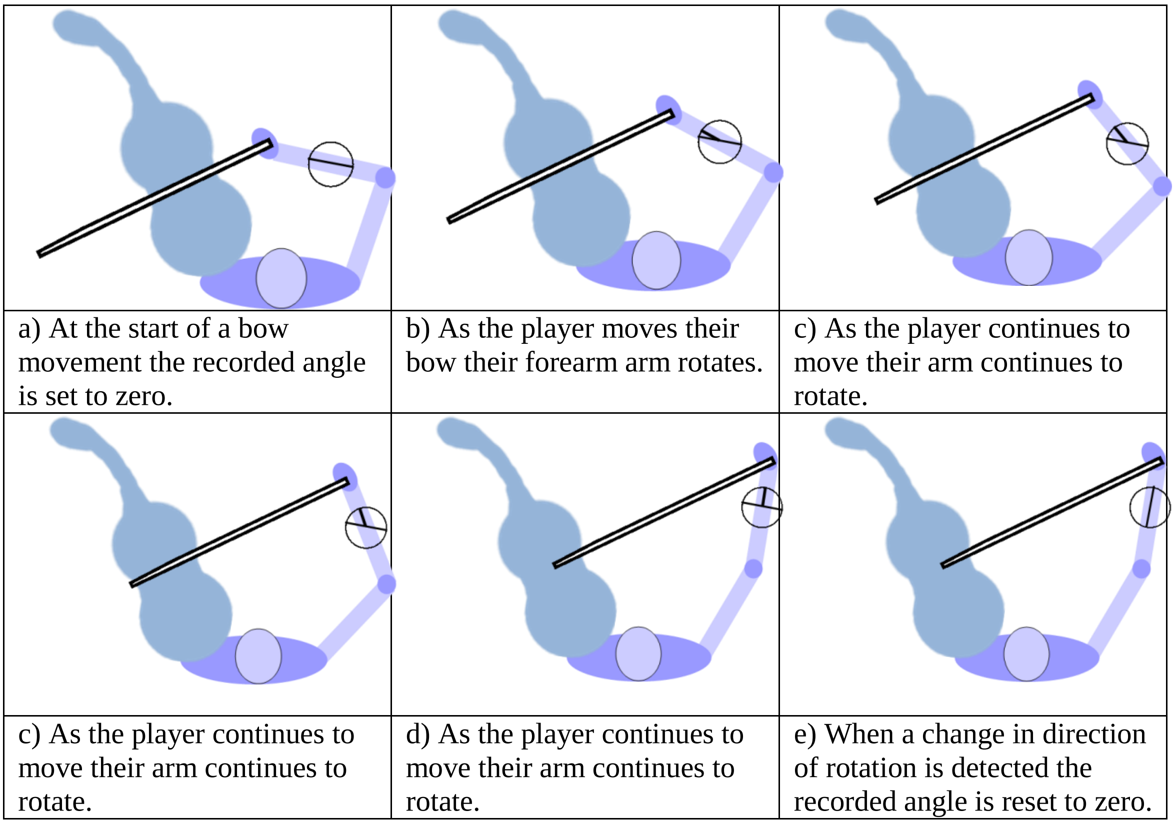

Figure 5-2: How the angle of the player's arm relates to the length of bow used

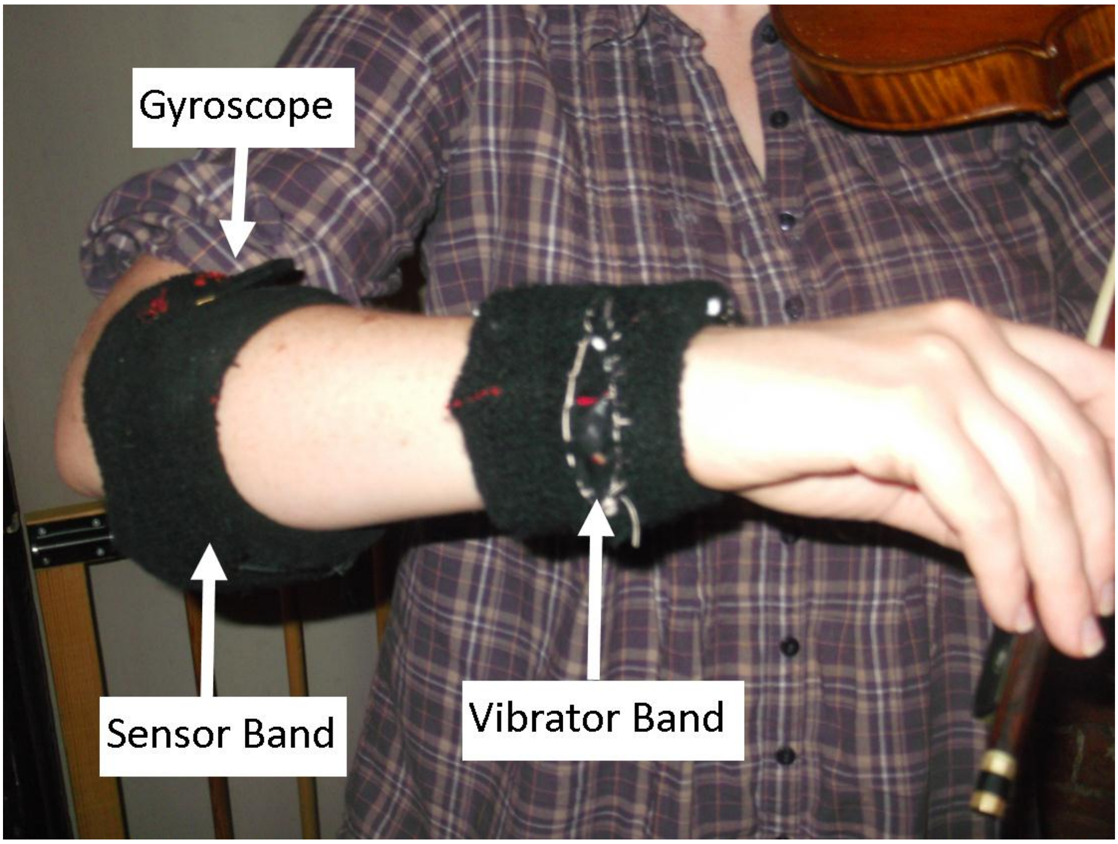

Figure 5-3: The sensor band worn on the forearm and the vibrotactile band worn on the wrist.

Vibrotactile Band

The vibrotactile band was another sweat band and is designed to be worn on the wrist (Figure 5-3). Inside the band there was a vibration motor (10mm shaftless DC motor from (Precision Microdrives, 2014)) controlled by an Arduino Mini Pro via a Darlington transistor pair. The vibrators were chosen for their small size, low cost and simplicity. They enable some control of vibration intensity based on the voltage used. However there were also some limitations to using these motors, they have a lag time between being switched on and the vibrator getting up to speed – this is exacerbated if less intense vibrations are used therefore the full intensity was always used (3.3V). They also do not allow more complex vibrotactile sensations such as rough or smooth vibration.

The Arduino was also connected to an XBee series 1 radio transceiver which received the measurements from the sensor band. Based on whether the measurements passed a certain threshold the Arduino switched the vibrator on or off. If the player had moved their arm through a certain angle, defined as being a 'long bow' the vibrator was switched on, the vibrator remained on until the player changes bow direction. If the player does not use enough bow so that their arm has not moved through the threshold angle they will not feel a vibration. The sensation when playing this feedback is that you feel a short buzz at each end of your bow when you are playing with lots of bow and you miss the buzz if as you shorten the amount of bow you are playing with.

The Visual Prototypes

Technologies such as easily portable screens and tablets could enable visual feedback to take the form of annotations on a digitised version of the sheet music. This has potential advantages in that it would not require the players to shift their gaze from the music in order to see the feedback. However, there are drawbacks to this approach as well. All music that participants might wish to play would have to be digitised and this would place limitations on how the feedback could be used. Participants would not be able to use their own music. Also depending on how the system is designed to follow the music being played, it may limit how easily players can replay difficult sections and still get feedback in the correct place on the music. In addition, combining the feedback with the musical notation does not ensure that it will be noticed. Hommel et al. (2001) describe in their theory of event coding that what we perceive is controlled by our intended actions. Thus, if a player is very focused on following the musical notation, he or she will still only notice the position of the notes on the stave, even though there may be additional annotations incorporated into the notation.

Based on this reasoning we chose a simpler, more flexible, approach to visual feedback. This took the form of a row of LEDs positioned on the music stand above the sheet music. This meant that participants could use their own sheet music, or use no sheet music at all, depending on how they would normally practise. This approach was well suited to the in-the-wild methodology which aims to study feedback being used for realistic learning tasks that are relevant to participants outside of the study setting.

A light bar consisting of 22 RGB LEDs in a line was designed to be positioned above the music on the music stand (see Figure 5-4). These LEDs are controlled in pairs by an Arduino Pro Mini via two TLC5940 PWM chips which were used to increase the number PWM outputs from the Arduino. The Arduino was also connected to an XBee series 1 radio transceiver which reads in the measurements from the Sensor Band. Based on these the Arduino would switch the LEDs on and off to display feedback to the player. The Light Bar has two feedback settings, one we refer to Simple Lights and the other as Complex Lights.

(i) Simple Lights condition

The simple lights setting is analogous to the vibrotactile feedback. The lights only come on when the player moves their bow far enough to register as being a full bow. When this happens the entire bar of LEDs above the music lights up green. When the player then changes their bow direction to start a new bow the lights switch back off. This setting was chosen so that it could be compared with the vibrotactile feedback, in that it communicates the same information about bowing in the same way but using the visual modality rather than the tactile modality.

(ii) Complex Lights condition

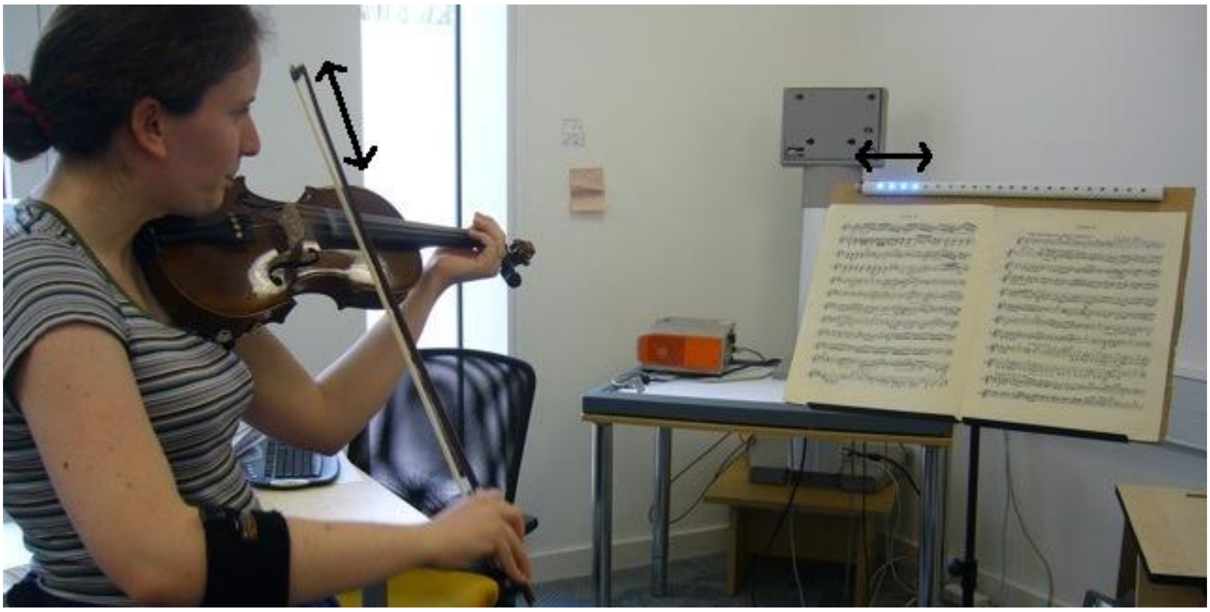

Figure 5-4: The complex lights gradually fill up as the bow moves across the string

The complex lights setting shows more information than the simple lights. It was designed to exploit the capacity of the visual modality to show a higher resolution of information than the tactile modality. The display shows the length of bow the player has used so far as they play. The LEDs come on gradually along the light bar in increments of two to represent how far along the bow the player currently is as they move it. At the point where a full bow would be registered the colour of the lights changes from blue to green. This setting was chosen to be representative of more complex feedback than could be provided using a tactile display. The reason for giving more complex feedback is that it could provide more clues as to how playing needed to change when practicing. The incremental display of the complex lights has two potential advantages over the single visual or tactile displays. First, by giving continuous feedback about the player's movement the link between movement and feedback may be easier to understand, making interpreting the feedback easier. Second, by showing the build up toward the full bow where the green lights come on it may motivate players to move their bow further since they can see the incremental progress towards their goal.

The display could have been designed to provide more information about the player's performance, by providing more patterns, colours, etc. However, this could have made it too complicated to be used in the ambient visual field and would interfere with the player's ability to read the music. Hence, the aim was to determine if providing further information in an incremental way enhanced practice or whether keeping feedback very simple is the best approach. By giving two levels of complexity of visual feedback we aimed to examine the trade-offs between two competing concepts in the theoretical framework – namely cognitive/perceptual load and motivation. On one hand, the simple feedback may be more effective because the flashing gives a clear signal to attract exogenous attention (Mulckhuyse & Theeuwes, 2010). On the other hand, the more complex moving display may be more effective because it shows incremental progress towards the long bowing goal which may be more motivating; moreover the continuous feedback will make a clearer link between movement and feedback.

Calibration

The sensor band requires a threshold to decide when a full bow has been used. In the study in this chapter we chose not to calibrate this threshold for each individual participant, instead we aimed to make the prototype one-size-fits-all. This was to save participants time as study sessions were often in the participants' lunch hour and could not over-run.

Finding a one-size-fits-all calibration

The choice of the threshold between a long bow and a short bow needed to be chosen carefully in order to meet the needs of the multiple participants. The sensor does not measure the bow length directly but rather the angle of the player's arm as they move the bow. Working with this kind of sensor meant that a threshold angle of the arm had to be chosen rather than a specific length along the bow. The exact length of bow this related to would then vary slightly based on the player's arm length and body shape. Initially the threshold was chosen based on the researcher's own playing so that the positive feedback was triggered approximately 10cm from the end of the bow. However, when initially trialled on this setting it was found to be very difficult to use as the participant had much longer arms than the researcher which meant he needed to play using the entire bow to trigger the feedback. Based on this participant's experience the threshold was changed so that the positive feedback was triggered more easily. Having found a compromise that worked for both longer and shorter armed people, this threshold was then kept constant for the study. This first test participant's bowing data was excluded from the analysis.

This first test case also revealed that the vibrotactile feedback came on later than the visual feedback. This is because there is a lag time for the motor to get up to speed which meant that it was felt further down the bow than the visual feedback. Therefore, the vibrotactile feedback was set to be triggered slightly before the full length was reached so that the motors were active by the time the threshold was reached. If a player did change bow direction after the vibration was triggered but before the full bow threshold was reached it would be unlikely the participant would feel the early vibration as the motor would not have had sufficient time to start moving before it was turned off again by the change of direction.

5.4 User Study

This user study was designed to investigate the differences between the visual and vibrotactile modalities in terms of their efficacy in delivering real-time feedback to people playing the violin. Both quantitative and qualitative data were collected. The quantitative data was collected to compare the efficacy of the feedback for improving bowing length and test the hypotheses. The qualitative data was collected to reveal how the modalities differ in communicating information to participants. It was also used to build a framework for understanding multimodal feedback.

To test the two hypotheses and compare the different modalities qualitatively, a within participants study was devised (Table 5-1). It consisted of individual practice sessions in which participants from an amateur orchestra tried every type of feedback. Each participant began their session by playing without the feedback, under two conditions: one where they practised normally and one where they focused on long bowing. Next they practised using each type of feedback in a counterbalanced order. Lastly, they practised again without the feedback. In each condition, the participants played the piece multiple times. Comparing quantitative bowing data from the feedback conditions will test the hypotheses concerning the effectiveness of different modalities for giving feedback. Qualitatively analysing a post-practice interview with each participant will allow us to study in depth the differences between the modalities.

Condition	Description	Order
Normal Bowing	Practice normally without feedback	First
Long Bowing	Practice focusing on long bowing without feedback	Second
Vibrotactile	Practice with the vibrotactile wristband to encourage long bowing	Feedback conditions given in a counterbalanced order for each participant
Simple Lights	Practice with the lightbar with the simple light display to encourage long bowing	
Complex Lights	Practice with the lightbar with the complex light display to encourage long bowing	
After	Practice focusing on long bowing	Last
The study concluded with a 10 minute interview about the feedback		

Table 5-1: Experimental design

The within participants design meant that each participant experienced the three different kinds of feedback: simple visual, complex visual and vibrotactile. With violin practice it is likely that participants would improve from their time spent practicing without the feedback especially once they knew that the intention was to increase bow length. This is why they were asked to play the pieces multiple times in each condition (approx. 3-4 times per two minute practice session) to see whether there were trends within each condition that suggested learning and whether the feedback conditions produced a sudden change. It was decided not to use a between participants design or have a control group because the studies with the MusicJacket (3.7) had shown strong individual differences in the way people learn and play the violin, moreover the participants were a mixed group with some very experienced players and some less experienced. This meant that it would be very difficult to draw comparisons between groups, particularly because the total number of participants was limited by the size of the orchestra.

5.4.1 Participants

Fifteen participants were recruited from a university orchestra who played in either the first or second violins, viola or the cello sections. They covered a wide range of ages (20-65+) and abilities and mainly consisted of university staff or their family or friends. All participants were known to the researcher as she was also a member of the orchestra. Participants were initially recruited using a general email to all the string players in the orchestra and many participants volunteered in response to this. In order to recruit more participants, the researcher also approach members of the string section individually either in person at rehearsals or by email. These individuals were selected based on recommendations from existing participants and those that the researcher thought might be particularly interested in participating. By the end of the study nearly all the regular first and second violin and viola players had participated, but only a small proportion of the cello section. At the start of the study participants were asked to sign a consent form saying that they agreed to take part but could choose withdraw at any point. At this point most of the study had been explained to them, but they accepted that they could not know exactly what the prototypes were measuring until after the first condition was complete. Participants were given £5 gift voucher for a store of their choice for taking part.

5.4.2 Procedure

The researcher met with each participant individually for one session which lasted approximately an hour. During this time they were asked to practise some sections from the music they were rehearsing for an up-coming concert. Studies with the MusicJacket had shown that feedback was more effective when it was relevant to a goal the participant was already trying to achieve in their playing (van der Linden, Johnson, Bird, Rogers, & Schoonderwaldt, 2011). To give this kind of relevance, the sections were chosen by the conductor as parts where the string section needed to maximise their bow usage. He had also already pointed out these sections to the strings as needing more bow in the natural course of previous rehearsals.

The practice sessions took a similar format each time (see Table 5-1). First participants were asked to play two particular sections of their orchestra music for two minutes each as if they were practicing it normally. The purpose was to get a baseline of how they play normally without any particular practice focus. The researcher then revealed to the participant that these parts had been selected by the conductor as parts where the full length of the bow needed to be used and that the focus of this study was on long bowing. This would make sense to the participants because these sections had already been pointed out during orchestra rehearsal as requiring this focus. Participants were then asked to play the pieces again for two minutes each with this practice focus in mind. This was in order to see how much improvement they were able to make by changing their practice focus. The participant was then introduced to one type of feedback. This was simple lights, complex lights or vibrotactile. The order of the feedback was changed for each participant so as to avoid order effects when comparing the different types of feedback. Before playing the pieces again the participant was encouraged to familiarise themselves with what the feedback did by playing a scale or other simple exercise. This also tested that the feedback was working in each case before measurements were collected. They then played the selected sections, again for two minutes trying to use the feedback to help them maximise bow length. The participant was introduced to the other two types of feedback in turn and asked to go through the same processes for each.

While the participants played, bowing data was collected in all conditions in order to compare the effectiveness of the different feedback types in promoting long bowing. The participant was asked to play through each section twice more without the feedback as a post feedback measurement. This was in order to see whether any improvements with the feedback were sustained when playing without the feedback or whether they returned to their initial bowing style when the feedback was removed. The sessions were held in practice rooms on campus, or when these were unavailable, a usability lab on campus.

Throughout these sessions, participants were encouraged to approach their practice as they normally would at home, so they could stop and start or replay sections if they chose. All that was specified was that they should practise it as they normally would but to try and keep playing for as much of the time as possible. Once the aim of longer bowing had been introduced it was reiterated on the introduction of each new type of feedback and again for the after measurement. The reasoning for not controlling how people practised the sections was three-fold. First, there is the question of ecological validity; controlling players to such an extent that every data point was comparable would remove all elements of natural practice from the study, removing the relevance of my findings for the overall goal of designing feedback for violin practice. Second, by implementing such control participants would not be able to experience the feedback on their own terms which would severely limit the discussion in the interviews. Third, from a pragmatic point of view, ensuring that each time the piece was played it was at the same speed and without pauses would be extremely difficult to achieve, particularly given the broad range of abilities.

5.4.3 Qualitative Data Collection

After using the feedback the participants were interviewed about their experience. This lasted about 10 minutes. The interview was semi-structured; the opening question asked the participants to rate the different kinds of feedback in preference order. Participants were then asked to explain why they had chosen that order and encouraged to talk about each one in turn. The questioning was kept broad and took the lead from the participant so that participants could talk about what was important to them about the feedback. Participants were then asked if they would use a similar system at home if it was available. Finally, they were asked what they would like a system to do if one was made just for them. They were encouraged to cover both what they would like the system to sense and how they would like to be given feedback. There were several reasons for asking these questions. One was to find out how relevant the feedback focus was in their eyes – whether they would normally focus on bowing or whether there was another element of playing they considered more of a problem. Another was to find out whether they envisaged a different type of feedback working better for them and why. It was also useful from a design point of view to hear what players would design for themselves. The interviews were audio recorded and then transcribed.

5.4.4 Quantitative Data Collection

The main measure used for the quantitative analysis of bowing was the percentage of long bows that were played by the participant each time they played the piece of music. A long bow was defined as a bow long enough to get positive feedback from the system. This corresponds to using more than three quarters (approx.) of the full length of bow, which was similar to the amount of bow mimed by the conductor when describing long bows in rehearsals and in the consultation prior to the study. Each time they played the piece the number of long bows (l) and short bows (s) was recorded and this was converted into a percentage (100 * l/(l+s)).

A method of data collection was chosen in order to have synchronised bowing data and audio data of discussion and playing so that the researcher could interpret the bowing data in the context of what each participant was playing at the time. Whilst playing, participants wore a sensor band on their arm which measured the length of bow they were using whilst playing in all conditions (normal practice, long bowing practice, vibrotactile, simple lights, complex lights and after feedback). For data collection, measurements were sent to the researcher's computer, the computer then flashed a red screen if the bow was too short to get positive feedback and a green screen if the bow was long enough to receive positive feedback from the system. Only the researcher could see this screen. The computer screen was recorded by a video camera which was also synchronously recording audio of the participant practising. The number of long bows and short bows was noted down by the researcher from this video and used to calculate the percentage of long bows used each time the piece was played.

How participants played the sections during the two minute blocks was not prescribed. They could play at different speeds and practise smaller sections. This has an effect on how the sensor data should be interpreted because many of the readings are highly dependent upon the context they come from particularly because the pieces they are playing are not homogenous; some sections are more difficult than others. The audio data allowed the researcher to identify different practice strategies participants used, for example, breaking the piece down into chunks, and also the comments made by the participants while using the feedback. In some cases, this audio was used to collect timing data about how quickly participants were playing, as this can also explain variation in bowing data.

5.5 Results

The findings were analysed quantitatively and qualitatively. To begin ANOVAs were performed on the bowing data (percentage long bows) to test the two hypotheses. To show the main effect between the different types of feedback the sequencing of the results from the feedback conditions was normalised.

5.5.1 Quantitative analysis

The raw data collected from the participant's bowing had been converted into the percentage of long bows out of the total bow strokes each time the piece was played. The definition of a long bow was a bow that would trigger positive feedback from the system - approximately ¾ of the overall length of the bow. The average percentage of long bows used was calculated for each participant in each condition by taking the mean of results from the multiple times the piece was played in each condition (Figure 5-6). This was combined into a mean over all the participants for each condition, which is displayed in Figure 5-5. As a comparison for interpreting the graphed data the ideal number of long bows to use in the piece of music has been labelled as approximately 60-70% because this was the number used by the leader of the orchestra.

A repeated measures ANOVA was performed on the data. There were no significant differences between the performance of participants in the difference feedback conditions $F(2, 20) = 0.14$, $p = 0.986$ (the data met the assumption of sphericity). Therefore both hypotheses were not supported. Vibrotactile feedback was not found to be more effective at encouraging long bowing than visual feedback ($p > 0.05$); and complex visual feedback was not found to be more effective than simple visual feedback ($p > 0.05$).

Extending the analysis to examine the data in the non-feedback conditions showed significant differences between some of the conditions. Due to the large number of conditions compared to the number of participants we chose to carry out a limited number of pairwise comparisons using paired t-tests and then manually apply a Bonferroni correction to the confidence interval used to define a significant result. Four comparisons were carried out, plus the ANOVA to test the hypotheses involved a further three comparisons, therefore we applied a Bonferroni correction corresponding to seven comparisons to our significance requirement of $p < 0.05$ making our new requirement for significance to be $p < 0.0071$. The results of the pair wise comparisons are given below.

Normal Practice and Long Bowing Practice

The percentage of long bows using in Normal Practice were compared with percentage of long bows used in the Long Bowing measurements using a paired t-test, in order to see whether asking participants to focus specifically on bowing has a significant effect on their playing. A significant improvement was found between these two measurements (Δ = 7.63 (% full bows), p = 0.006). Δ indicates the difference between the two means.

Long Bowing Practice and Practice with Feedback

The percentage of long bows used in the Long Bowing condition was compared with percentage of long bows used in the feedback conditions using paired t-tests, in order to see if feedback could offer significant additional benefits to simply focusing on long bowing. A significant improvement was found between Long Bowing and all the feedback conditions: Vibrotactile (Δ = 8.98 (% full bows), p = 0.005), Simple Lights (Δ = 9.19 (% full bows), p = 0.001), Complex Lights (Δ = 8.78 (% full bows), p = 0.007).

These results suggest that all three types of feedback improved participants' bow usage. However, the sequencing between the non-feedback conditions and the feedback conditions was not normalised so we cannot discount that some of the improvement may be due to the effect of repeatedly playing the same piece. To look at this finding in more detail we examine the time order of the feedback on their performance (Figure 5-7 and Figure 5-8). One might expect that a training effect would show the same gradual improvement across all ordering. We can see from Figure 5-7 that although there are large improvements over the first three conditions there is no general upward trend across all conditions that would be indicative of gradual improvement due to repeated playing. This suggests that although training effects will have occurred, it is likely that the feedback is having an effect beyond this.

To conclude, the statistical evidence does not support the first two hypotheses about there being differences between the feedback conditions on performance. However, the individual bowing data does show strong individual differences between the feedback types. To further understand how different types of feedback may be more effective for different individuals we now give a qualitative analysis of the responses of each individual to the feedback, categorising them in terms of how the feedback changed their playing.

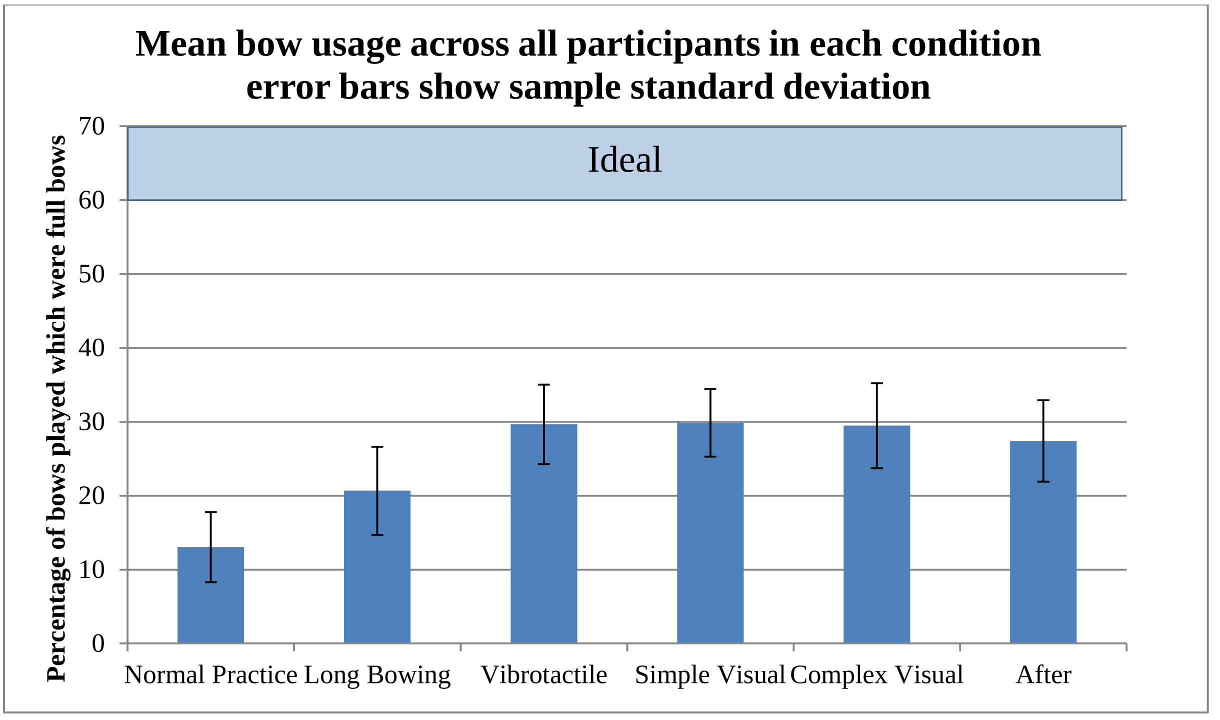

Figure 5-5: Mean bow usage across participants in each condition. Bow usage is measured as the proportion of bow strokes played by the participants classified as being full bows. For the piece being played 60-70% full bows is an appropriate goal, based on the way the piece was played by the lead violinist of the orchestra. The feedback data here is not time-ordered as participants experienced the different types of feedback in different orders.

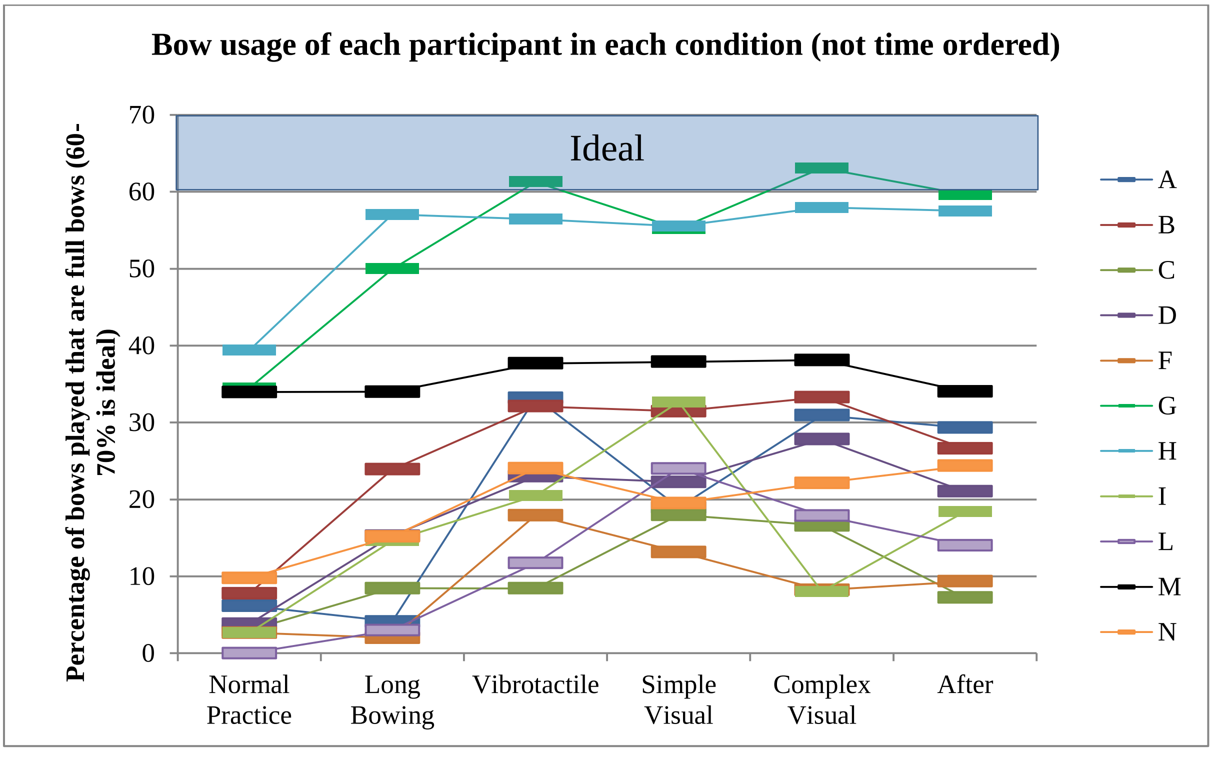

Figure 5-6: Individual bow usage data. Bow usage is measured as the proportion of bow strokes played by the participants that were classified by the sensor as being full bows. The data is not time ordered. The linking lines are added as visual aid to help keep track of participant data between different conditions, they are not intended to be meaningful in themselves.

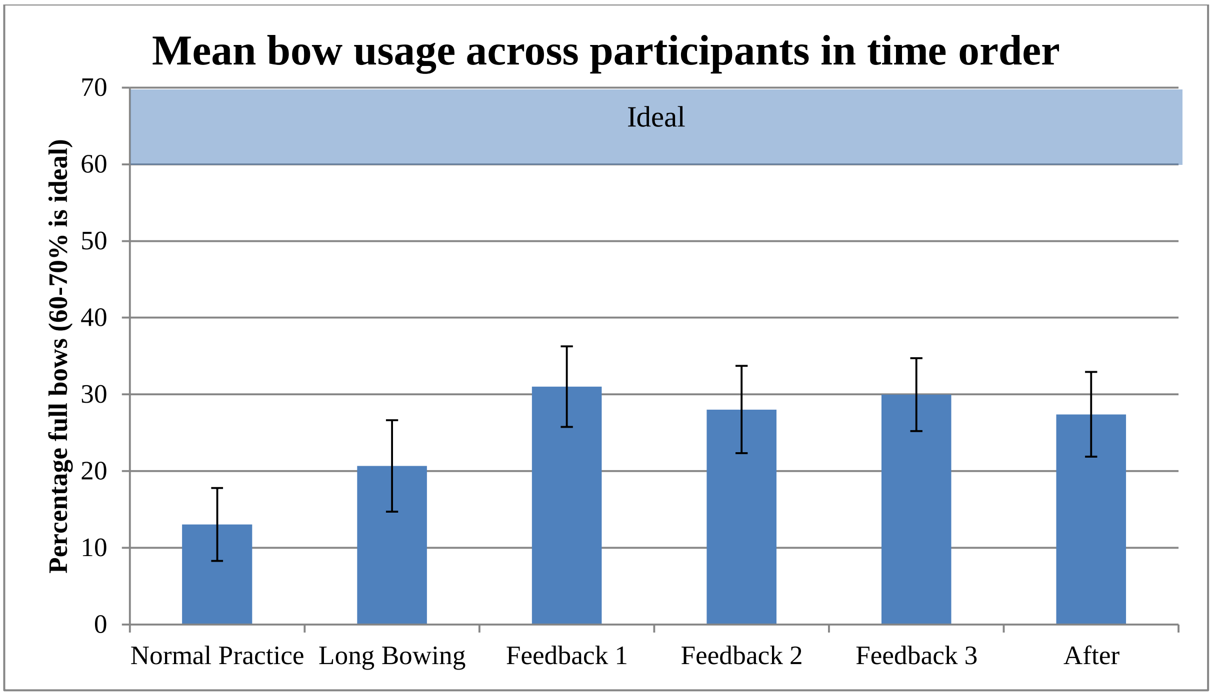

Figure 5-7: Mean bow usage across participants with the feedback grouped in the order participants experienced it rather than the type of feedback.

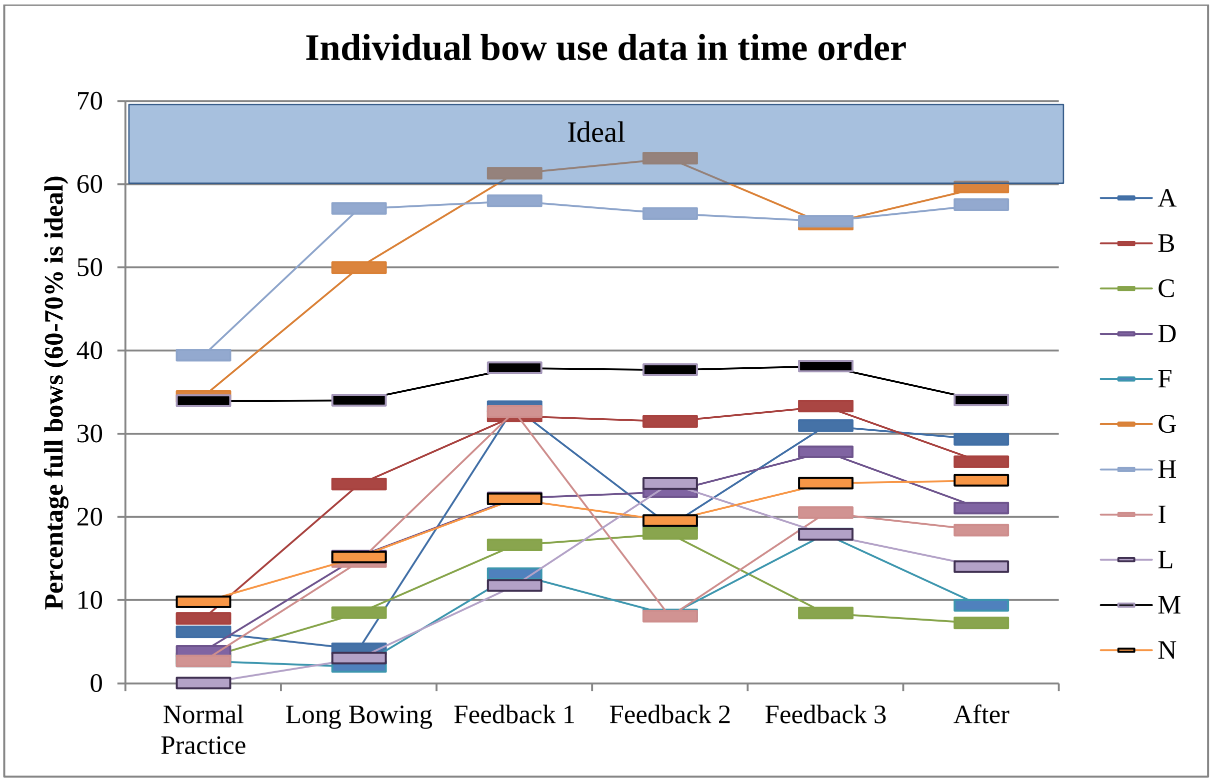

Figure 5-8: Individual bow use data shown with the feedback grouped by the order it was experienced in rather than particular type.

5.5.2 Qualitative Analysis

The bowing data and interview data were analysed qualitatively to examine individual differences in more detail. The bowing data was graphed for each participant and this was used to categorise them in terms of how the feedback influenced their playing. The transcripts from the interviews were then used to compare their experience and opinions about using the feedback with the outcome for their bowing. Extra information from the audio recordings of the practice was also drawn in to account for some of the patterns in the data. The participants are categorised in terms of different patterns of improvement (and no improvement). In the next section the data from each individual is analysed with respect to these five categories.

(i) Participants that only improved when using the feedback and did not improve without it

(ii) Participants that improved without the feedback and made further improvements, but only with specific types of feedback

(iii) Participants that show a general pattern of improvement over all conditions, with no clear link to the feedback.

(iv) Participants that adapted their practice strategy in the feedback conditions making improvement difficult to measure

(v) One participant showing no improvement over the session

Participant	Gender	Instrument	Favourite	Least Favourite	Category	Included in stats?
A	male	violin	Vibrotactile	Complex Lights	(i)Improves - feedback only	Yes
B	female	violin	Complex Lights	Simple Lights	(iii)General improvement	Yes
C	male	violin	Vibrotactile	Simple Lights	(ii)Improves - specific types of feedback	Yes
D	female	violin	Complex Lights	Vibrotactile	(ii)Improves - specific types of feedback	Yes
E	female	violin	Simple Lights	Vibrotactile	(iv)Changes practice strategy	No – practice inconsistent
F	male	cello	Vibrotactile	Complex Lights	(i)Improves - feedback only	Yes
G	male	viola	Vibrotactile	Simple Lights	(ii)Improves - specific types of feedback	Yes
H	male	violin	Complex Lights	Joint second	(v)No improvement	Yes
I	female	cello	Vibrotactile	Complex Lights	(ii)Improves - specific types of feedback	Yes
J	male	violin	Complex Lights	Vibrotactile	Additional interview data (no measurement)	No - data collection fault
K	female	violin	Simple Lights	Vibrotactile	(ii)Improves - specific types of feedback	No - data incomplete
L	male	violin	Vibrotactile	Complex Lights	(ii)Improves - specific types of feedback	Yes
M	male	violin	Vibrotactile	Complex Lights	(iii)General improvement	Yes
N	female	violin	Complex Lights	Vibrotactile	(ii)Improves - specific types of feedback	Yes
O	male	viola	Complex Lights	Vibrotactile	Additional interview data (no measurement)	No - incorrect calibration

Table 5-2: Participants who took part in the study and the categories they fall into

(i) Improvement with feedback only

Participants A and F are the most marked cases of showing improvement as a result of the feedback. Their data are particularly convincing because they were not able to improve their bowing between the normal practice condition and the long bowing condition where they were explicitly asked to focus on it as a goal without any feedback. This provides further support that their improvement was a result of the feedback and not due to a training effect of repeated playing because that hypothesis would predict a gradual improvement over these two conditions which is not reflected in the data. On being introduced to the feedback, their bowing performance was greatly improved as can be seen from Figure 5-9, Figure 5-10, Figure 5-11 and Figure 5-12.

Looking at their data in more detail reveals that the reason for this pattern is different in each case. Timing data taken from the audio recording for participant A shows that he increased the speed (or tempo) at which he played the piece in the long bowing condition instead of increasing his bow length. This response could be accounted for by understanding that in order to use more bow and play at the same tempo a player must move their bow faster. Here we can see an intention to use more bow by moving it faster but that it has not been carried through fully because he is still changing direction too early meaning that he plays the piece at a faster tempo rather than with longer bows. When the feedback is introduced (vibrotactile feedback) the tempo he plays at returns to a similar level to his normal practice but his bow length has increased suggesting that the external stimulus has encouraged him to reach further whilst still moving the bow faster than in his normal practice which results in playing the piece at the original tempo but with longer bows.

The reason for the improvement in participant F (Figure 5-11 and Figure 5-12) is different. After playing through the piece three times using the feedback (simple lights) he comments that *"it's a real effort actually to do that [play with a full bow]"* and out loud weighs up the pros and cons of using full bows in this case: *"I probably wouldn't off my own back seek to use very long bows on the dotted bits [near the start] because it's much more of an accented rhythm… but then you sacrifice the sound, in a professional orchestra they probably would be using lots of bow".* For the purposes of the study, he continues trying to play with long bows and deliberately chooses to play the dotted section near the start because it is the most difficult part to play with long bows. It is likely that his increased use of bow with the feedback comes in part from this new resolve to really aim to use the full extent of the bow and get the good feedback. Therefore, in this case the feedback and the study setting is actively encouraging more bow usage rather than simply enabling it.

The data also show that for both participants, the different types of feedback are successful to different extents in how many full bows they encourage. For both participants the vibrotactile feedback appears to be the most successful. This is also the feedback they both chose as their favourite in the interviews.

The reasoning participant A gave for his choice related to being aware of the feedback. He felt the complex lights were *"too sophisticated just for the sheer practicality that you can't really see the green"*. He added that with this display, he had trouble noticing it at all when playing music near the bottom of the page (the light bar is fixed at the top of the music stand) and that even when the music was at the top *"he couldn't really distinguish between whether it was two green or three green or four green"*. He had more difficulty giving a reason for choosing the vibrotactile over the simple lights saying *"there's not much between the two"* but he was clear when talking about future systems that could be designed for him that we had *"honed in on the buzzer [vibration]"* meaning that the he believed the vibrotactile was the type of feedback he would like to use in the future. The bow usage data shows a contrast in his playing between the vibrotactile feedback, the most successful type of feedback, and the simple lights, the least successful type of feedback. This shows that it is difficult for participants to evaluate how they are using a display, as in the interview he considered his awareness of the two displays to be equal.

Participant F also said that *"there's not a huge amount in it"* when it came to comparing the vibrotactile and the simple lights but that he preferred the vibrotactile because it was a different modality to reading music:

"because it's not something visual you've got to – it giving you a signal in a different way I think that is rather good. The lights are sort of more straight forward I suppose but er – straight forward I don't know how to describe…!"

Participant F is a cello player, and also initially commented that he had difficulty distinguishing the vibrations of the feedback from those coming from the cello but was able to overcome this quite quickly:

"so I think perhaps because I was playing quite loudly and that's quite a robust passage that first one, quite difficult to – sense. But I think then again once I got used to it then I started picking it up – think that's probably the reality. "

He discounted the complex lights because he found it more difficult to get the full bow feedback with them. This is also apparent in the bow usage data which shows less long bows measured. It is interesting that he should be aware that he is not able to get the good feedback but unable to do something about it, when he was able to in the other conditions and shows the importance of not simply communicating information but doing it in a way that encourages the player to take action.

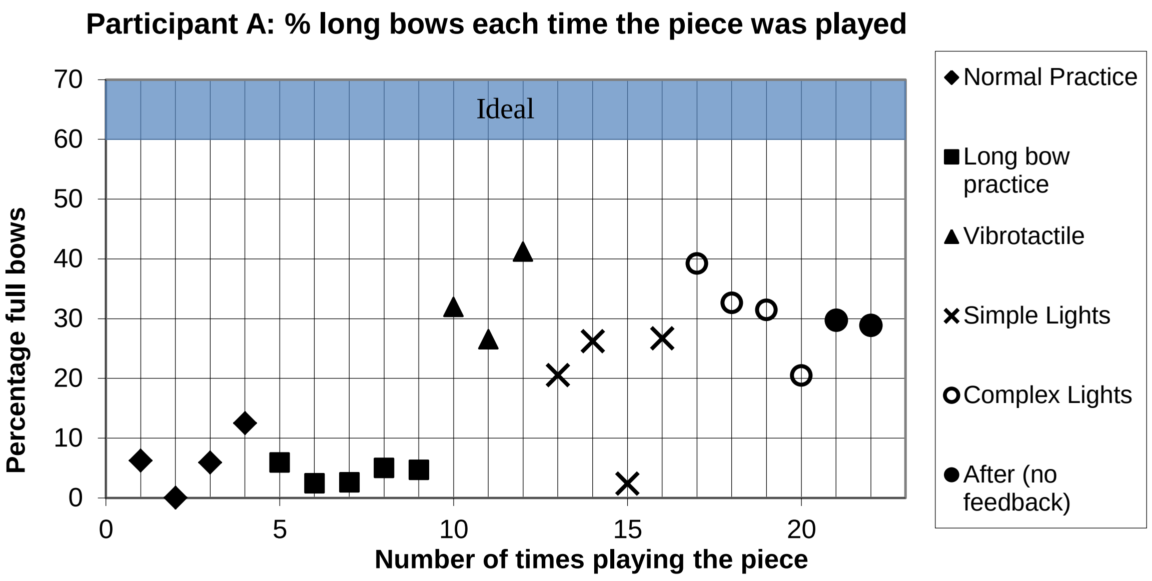

Figure 5-9: Percentage of bows played by participant A that were counted as full bows (60-70% full bows is considered an appropriate ideal for this piece). Each data point represents playing through the piece once apart from points 5 and 19 where he cut the piece short two thirds of the way through. All results are shown in the order that the feedback was given.

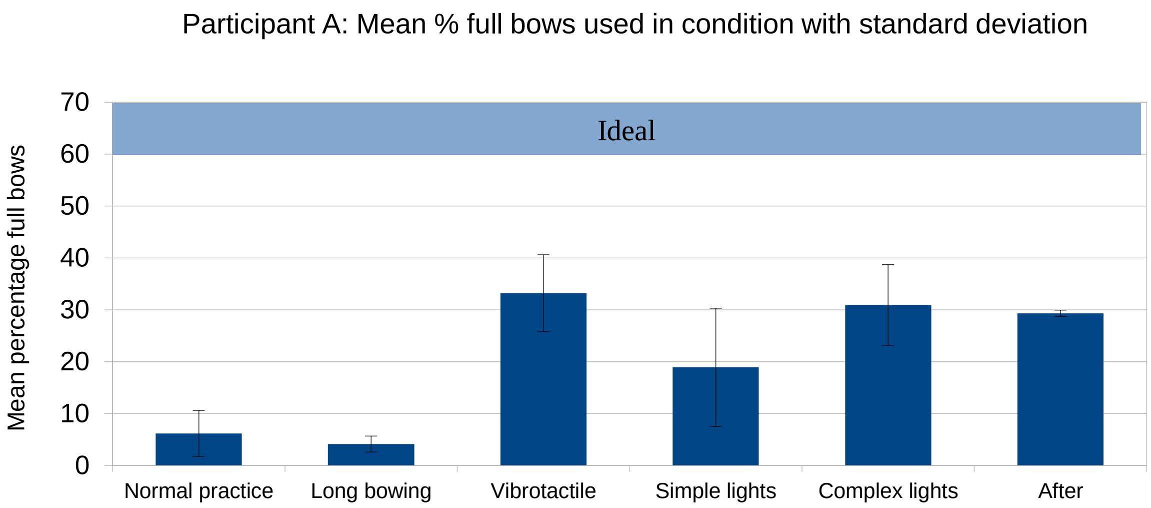

Figure 5-10: Mean percentage of full bows used by participant A in each condition (60-70% is considered ideal). This is calculated by finding the mean from the repeated times the piece was played in each condition. The error bars show the standard deviation.

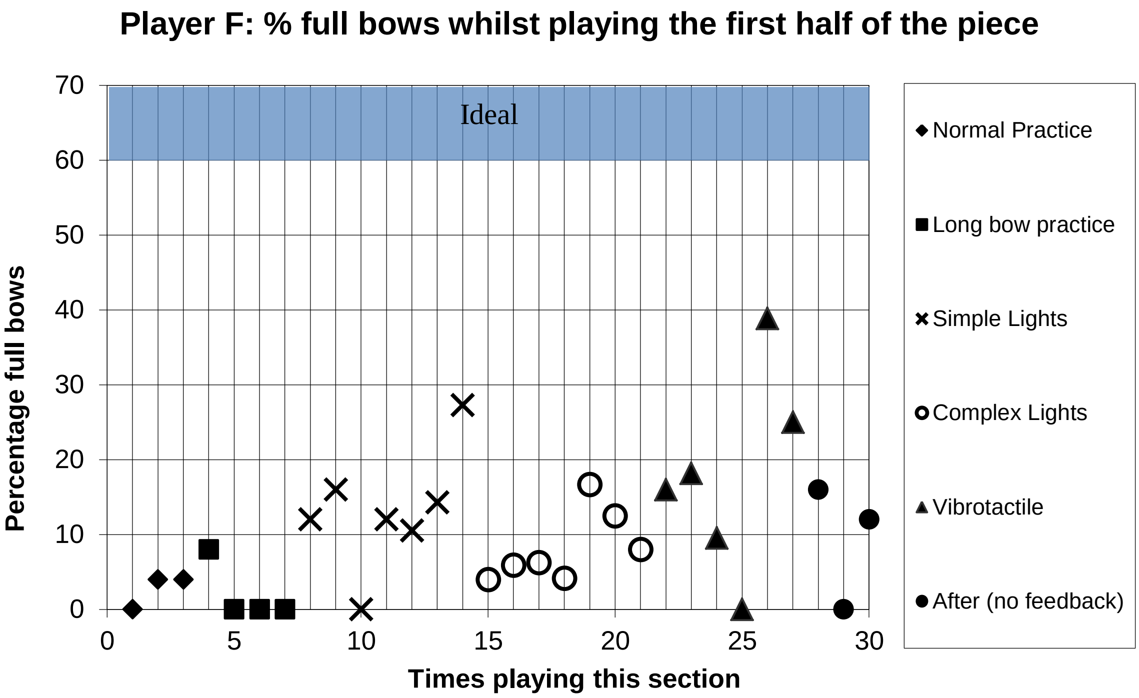

Figure 5-11: Percentage full bows used by participant F while playing the first 25 bows in the piece. Each point represents playing these 25 bows. This has been done because on getting the feedback, he chose to focus only on this particular part of the piece for all subsequent playing; therefore in order to make the points more comparable only the data from the first 25 bows in the piece have been used for all data points whether or not he then proceeded to play the rest of the piece.

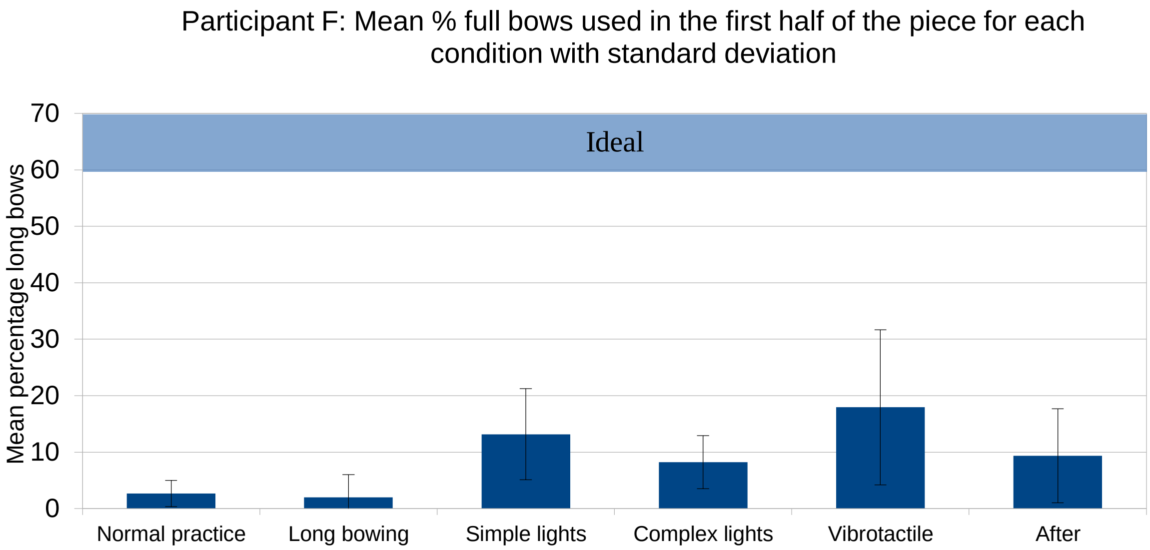

Figure 5-12: Mean percentage of full bows used by participant F for each condition. The error bars show the standard deviation. This graph shows a clear improvement between the initial conditions without the feedback and the introduction of the feedback.

(ii) Improvement with specific types of feedback

A further seven participants show clear improvements with the feedback. Like the two participants already discussed, they show larger improvements with specific types of feedback and less improvement with others (in some cases no improvement with certain types of feedback). This contrast between the different feedback conditions further contradicts the suggestion that the main effect is due to a training effect, since that case would predict a general improvement over all conditions whereas the data shows better performance only in specific conditions. Specifically, which types of feedback worked best is different for different participants. Participants were not always aware of which feedback was having the strongest effect on their playing. The following presents the results from each participant in turn with additional details added from the interviews and audio about their experience using the feedback.

Participant I showed the largest differences between effectiveness of the three types of feedback (Figure 5-13 and Figure 5-14). The data shows that for her the simple visual feedback was most effective followed by the vibrotactile and then the complex visual feedback. She was initially excited by the concept of the complex (moving) lights when she first saw them, however, after playing the pieces with them her opinion changed. In the interview she commented:

"Although I liked the moving lights when I wasn't trying to concentrate on the music I found it - not distracting - but I wasn't really concentrating on it because I was looking at the music so I couldn't really work out what it was trying to tell me."

This difficulty with seeing and understanding the feedback seems to be reflected in her bowing results where the complex feedback scored lower than when she was just focused on long bowing without any feedback. This may suggest that these difficulties are increasing perceptual demands or working memory load, which makes it harder for her to focus on the task of long bowing.

The vibrotactile feedback was her favourite type of feedback because she felt it was easier to understand:

"I didn't find it distracting and um it's kind of more relevant you're moving your arm and to have something happening in your wrist that – it makes you – it's just a bit more – my simple brain makes that connection more easily – 'oh yes I've got do something with my arm!'"

However, this was not the feedback that scored highest on the bow usage data. This suggests that she may have been too comfortable with the vibrotactile feedback and that it did not raise her awareness enough.

This concern is even more evident in the data of participant C (Figure 5-15 and Figure 5-16). His data shows that with both visual systems he performed better than the condition where he was concentrating on long bowing with no feedback. However, with the vibrotactile feedback, his bowing returned to similar levels as when he had no feedback. This contrasts strongly with the preferences he gave in the interview, where he chose the vibrotactile feedback as his favourite because he felt it appealed to a different part of his brain which made it less intrusive:

"*I think it's in a way less intrusive and it sort of fits in with the process of playing the violin. Um… to have a visual stimulus when you've got the music to play anyway is one added thing. I think the vibrations affect the cerebellum rather than whatever bits of your brain are taking the information from the lights."*

Later on, he describes the lights as giving a sense of *"doing it right"* but the vibration as being more like *"knowing"*. The contrast in bow usage between the visual and tactile conditions supports the idea that for him there is some clear differences in his experience with these two modalities; however, it suggests that the outcome is that visual feedback is more effective for lengthening bow use. It may be that the vibrotactile feedback is too unobtrusive and thus does not make him focus on bow length in the same way as the visual feedback which he felt required more conscious attention.

Participant L also shows a similar pattern. In the interview he was negative about the visual feedback, saying it was *"distracting"* and *"was taking my eye from the music"*. In particular he disliked the complex visual feedback because it was too difficult for him to process while reading music:

"It's a constant movement and I couldn't really get how – it's just overload for me […] trying to get an idea of what […] the blue and green lights were telling me was really too much".

This difficulty in interpreting the complex lights is reflected in the data as the simple lights were more effective for lengthening bow than the complex lights. He chose the vibrations as his favourite type of feedback because it was *"the least off-putting of the three"*. This is different to the reason participants I and C gave who felt it was easier to interpret as well as being less intrusive. Looking at the data from his practice (Figure 5-17 and Figure 5-18) shows that the vibrotactile feedback was the least effective for lengthening bowing. This shows that although participant L did not like being distracted, catching attention was a positive attribute for changing behaviour so long as he could interpret the feedback.

Participant K also felt that the lights caught her attention more than the vibrotactile feedback, but she saw this as a good characteristic. She rated the vibrations lowest because she felt they were *"easier to ignore"* and *"seemed quite constant"* suggesting she had difficulty perceiving when it was switching on and off. She also gave another reason for liking the lights which was unique among the participants. That is she specifically spoke about its helpfulness in enabling her to connect the feedback to where she was in the written score:

“whereas a light comes up and maybe you've got a bowing sign and it's at the end of the bowing sign and you're thinking OK I really need to make that bowing phrase longer, do you know what I mean, with your bow. Whereas the vibrations didn't seem to have that same... I wasn't able to as quickly associate that with a part of the music I think that's probably what I'd say. I'm only vaguely noticing it and I'm not noticing precisely where it fits in...”

This shows a different way of using the feedback; rather than connecting it directly to movement she is using the music as a way of structuring her interpretation of the feedback. In contrast she found the vibrotactile feedback difficult to use because she found connecting it to the music difficult. The complex lights she placed in the middle because she felt it was difficult to be aware of the more complicated display in her peripheral vision:

“I was not registering necessarily whether it was blue or green, you know what I mean? It was just that there was a light going. So I think that it was more helpful to me just to know light is right,”

Her bowing data (Figure 5-19 and Figure 5-20) shows that the vibrotactile feedback was least effective of all feedback types which fits with her perception that the vibrations were having least *“impact”* on her playing. Both types of visual feedback were equally effective in lengthening her bowing.

In contrast to the previous four participants, participant N found the vibrations more intrusive than the visual feedback:

“I think with the lights you can concentrate on other parts while you’re practicing as well and it can be more kind of subconscious whereas the only thing you can think of with the vibrations is that.”

This made her rate the vibrations as her least favourite. However, the bowing data (Figure 5-21 and Figure 5-22) shows that they were the most effective form of feedback. Again this shows that obtrusiveness, although having drawbacks, can make feedback more effective. It also shows that what modality of feedback is best at catching a participant’s attention is dependent upon the individual. She preferred the complex lights over the simple ones because it showed when she was *“nearly there but not quite”*. This extra incentive appears to have been effective as she did more full bows with complex lights than the simple lights.

These five cases show how the efficacy of feedback can be reduced when it is too *unobtrusive*. There were also cases where participants found the opposite, when feedback was too *intrusive*. One case is participant G who did not like the simple lights because he found the “flashing” of the lights was “too much” and he was not able to link the feedback to his movement easily. This is reflected in his bow use data (Figure 5-23 and Figure 5-24). This shows improvement between the normal bowing condition and the condition where he was focused on long bowing, it then shows a further increase in bow usage for the feedback conditions where he had either the vibrotactile feedback or the complex visual feedback. However, in the condition with the simple lights his bow usage went back down and although it was still better than the initial conditions without the feedback, it was worse than the after condition where he was also playing without the feedback. This suggests that the perceptual distraction caused by the “flashing” of the simple lights, coupled with the memory load of trying to decipher their meaning (which he commented he found difficult to *“link”* with his movement) adversely affected the performance of this participant.

The other two styles of feedback (vibrotactile and complex lights) were successful in improving performance. Of the two types of feedback, participant G preferred the vibrotactile because it was *“least intrusive, playing the music”* and he was aware of it *“as part of the action”* of playing violin. This description of the feedback is similar to the experiences described by participants C and I. However, the outcome is different as the vibrotactile feedback was almost equally effective as the complex lights which was the most successful type of feedback for this participant.

Participant D (Figure 5-25 and Figure 5-26) also shows most bow usage with the complex lights. Her data shows that she used more full bows in all conditions with the feedback compared to those without the feedback, but that she performed best with the complex lights. She chose complex as her favourite because it has an “extra bit” which made it more “definite”. She chose the vibrotactile as her least favourite because she disliked the sensation, describing it as “weird” and “disturbing”. This strong dislike is not reflected in the bow usage data as she performed equally well with vibrotactile as with the simple lights, and shows the importance of listening to people’s experience as well as studying effectiveness as it would clearly be inappropriate to use vibrotactile feedback in this case.

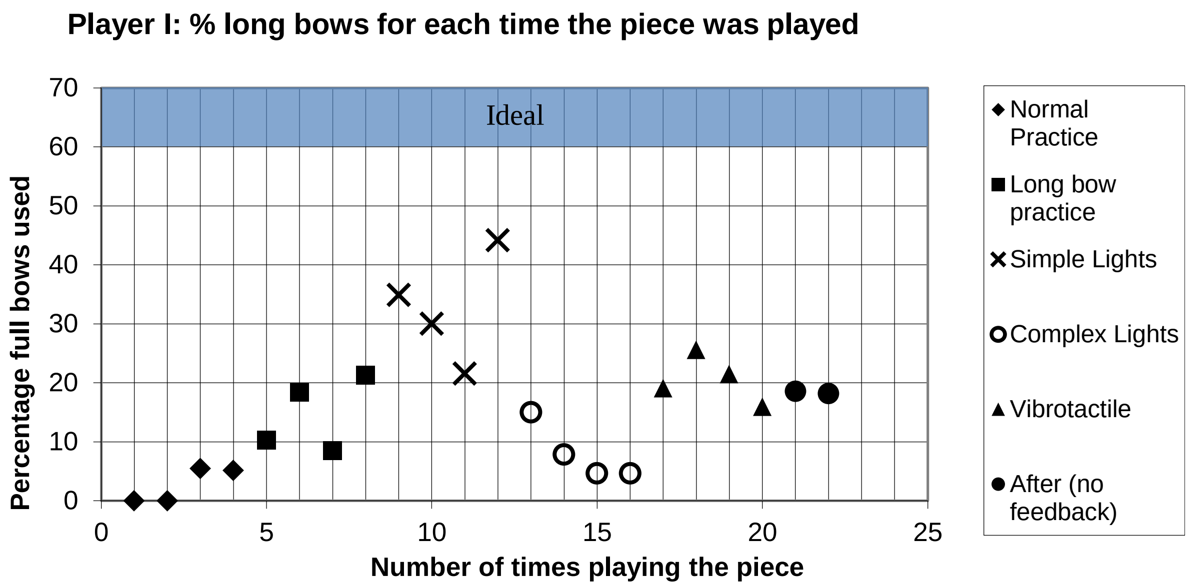

Figure 5-13: Bowing data from participant I showing the proportion of bows played that were classed as full bows, each time she played the piece. Each point represents one play-through of the piece and shown in the order she received the feedback. The data shows clear differences between the percentage of full bows she used with each feedback types.

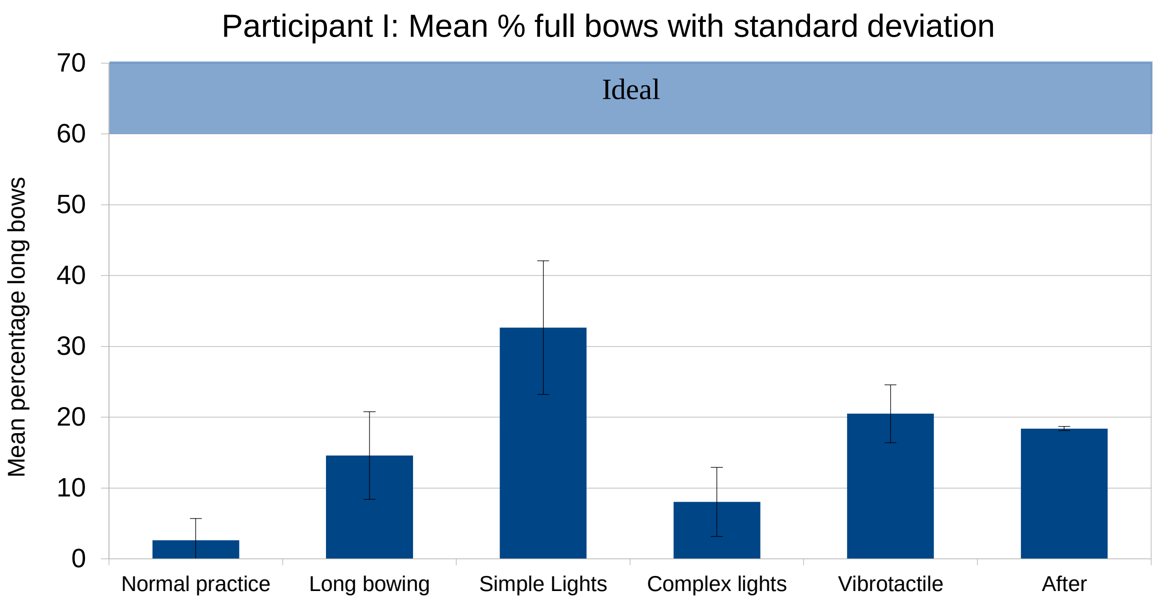

Figure 5-14: Participant I's mean bow usage in each condition, calculated by finding the mean across the multiple times she played the piece in each condition. Notice the clear differences in her bow usage with each type of feedback.

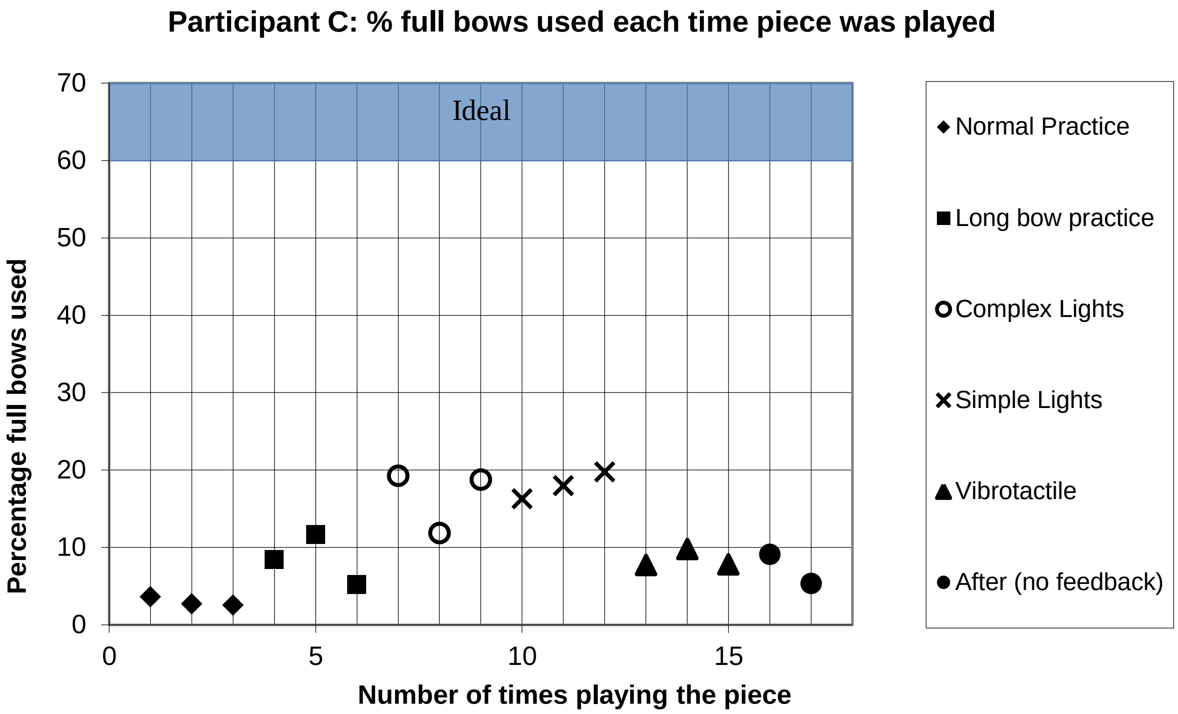

Figure 5-15: This shows the percentage of bows played by participant C that were full bows each time he played the piece. Notice that when the participant is given the vibrotactile feedback he returns to a similar way of playing as when he had no feedback.

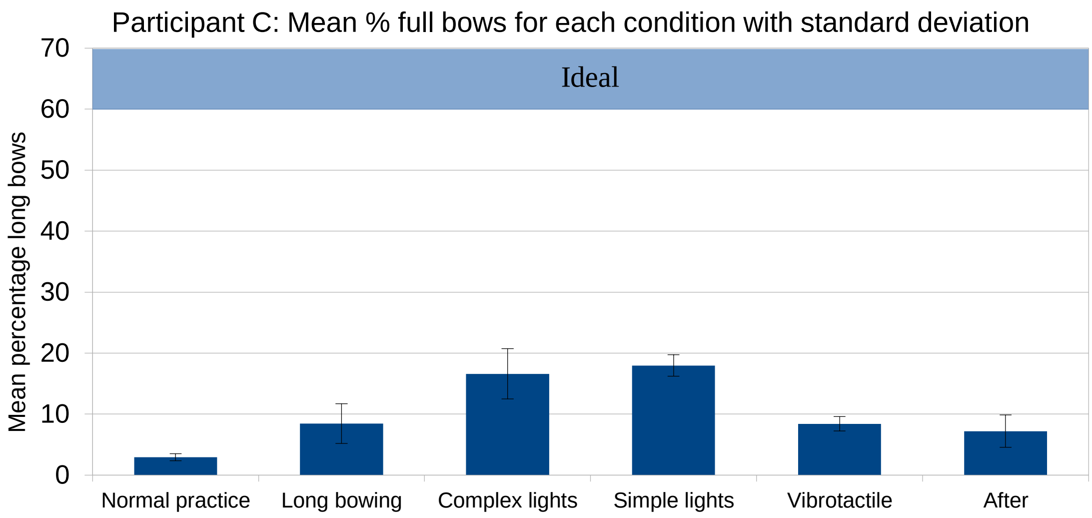

Figure 5-16: Participant C's mean bow usage in each condition, calculated by finding the mean across the multiple times he played the piece in each condition. The error bars show the sample standard deviation. Notice the clear drop in full bow usage when the vibrotactile feedback is introduced, even though this was his favourite form of feedback.

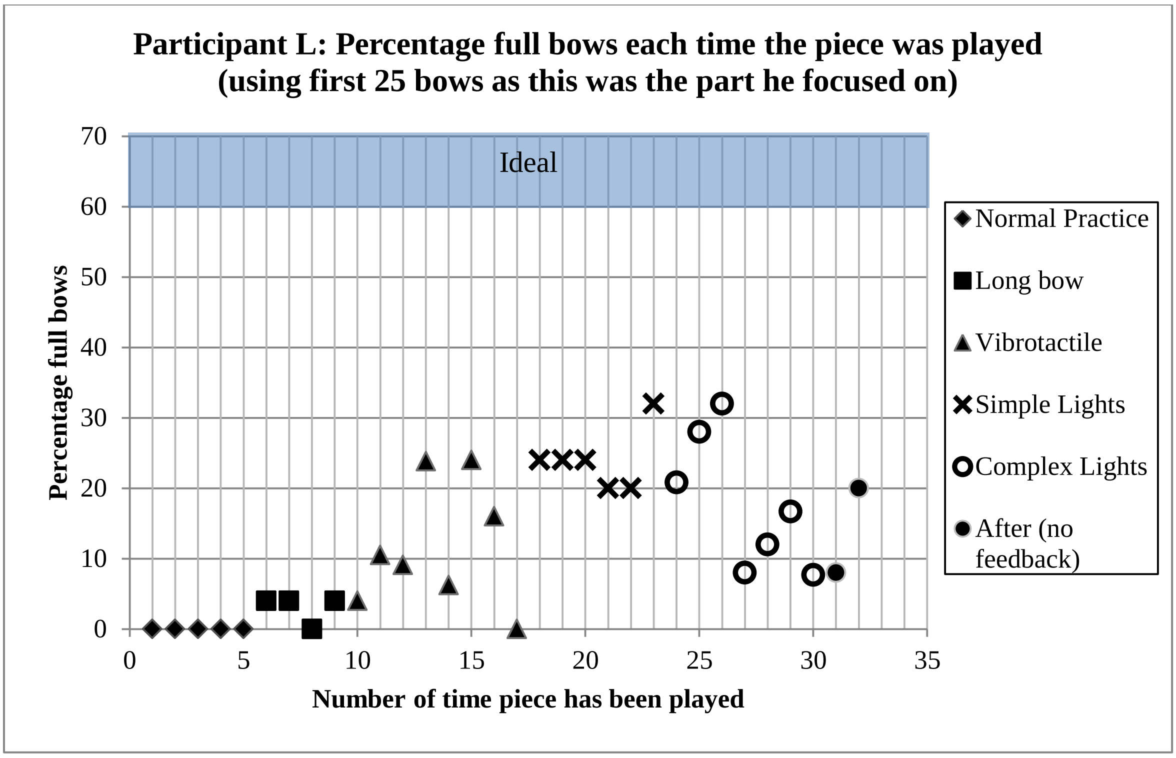

Figure 5-17: Percentage full bows used by participant L in each time he played the piece. Measurements are taken from the first 25 bows each time he played the piece as he focused his practice near the start and often did not play the piece fully.

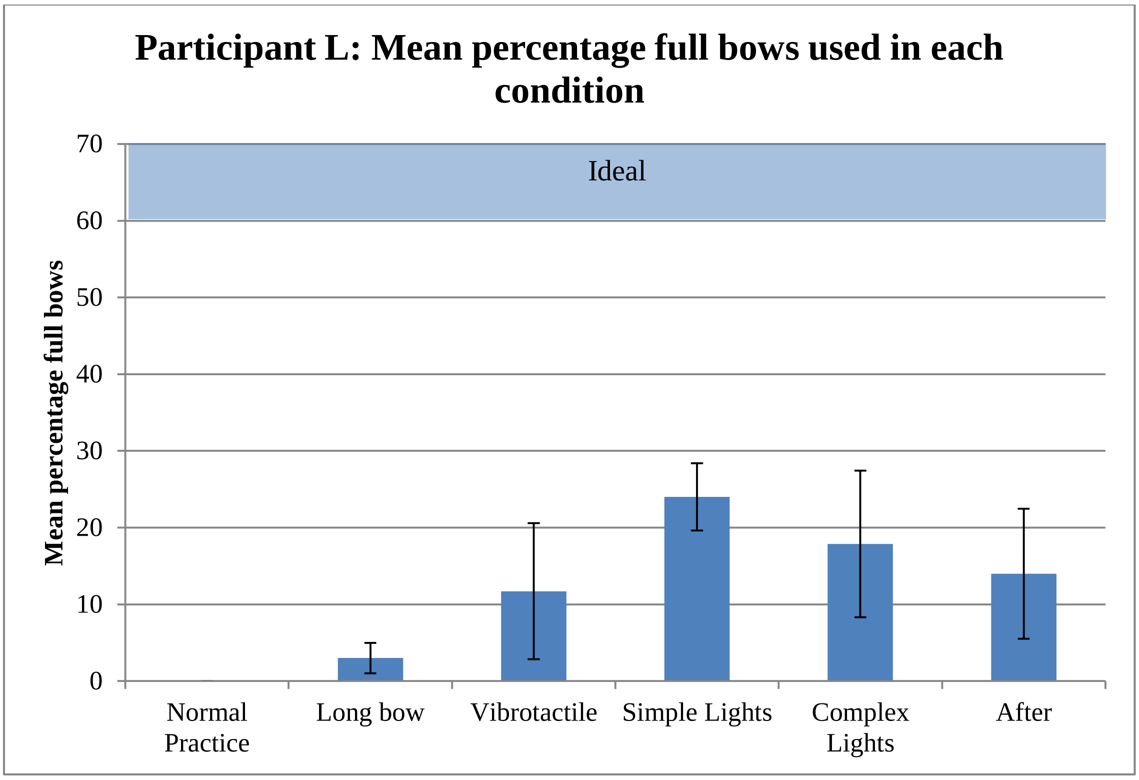

Figure 5-18: Mean percentage full bows used by participant L in each condition. Measurements are taken from the first 25 bows each time he played the piece as he focused his practice near the start and often did not play the piece fully. The error bars show the sample standard deviation.

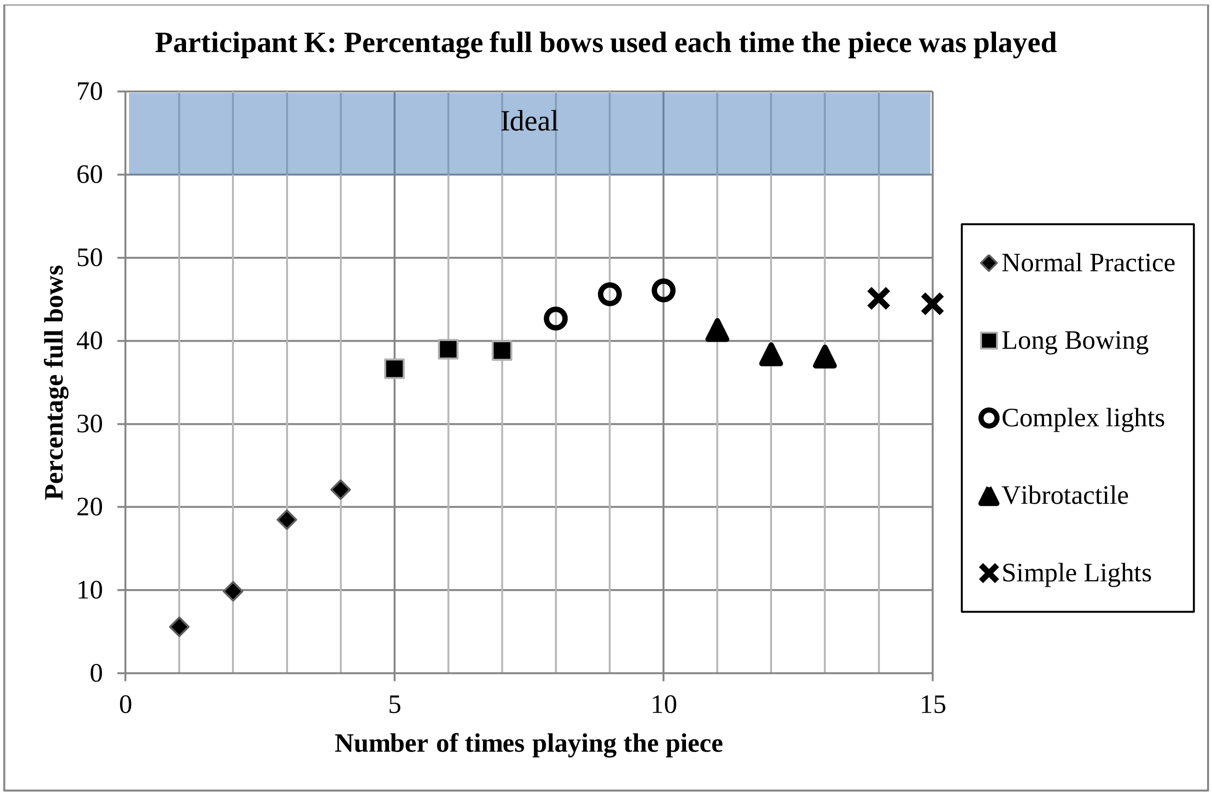

Figure 5-19: Percentage full bows used by participant K each time she played the piece. There is no after measurement in this data set due to a mistake in the recording process. Notice the dip in bow usage with the vibrotactile feedback, K commented that the vibrations did not attract her attention.

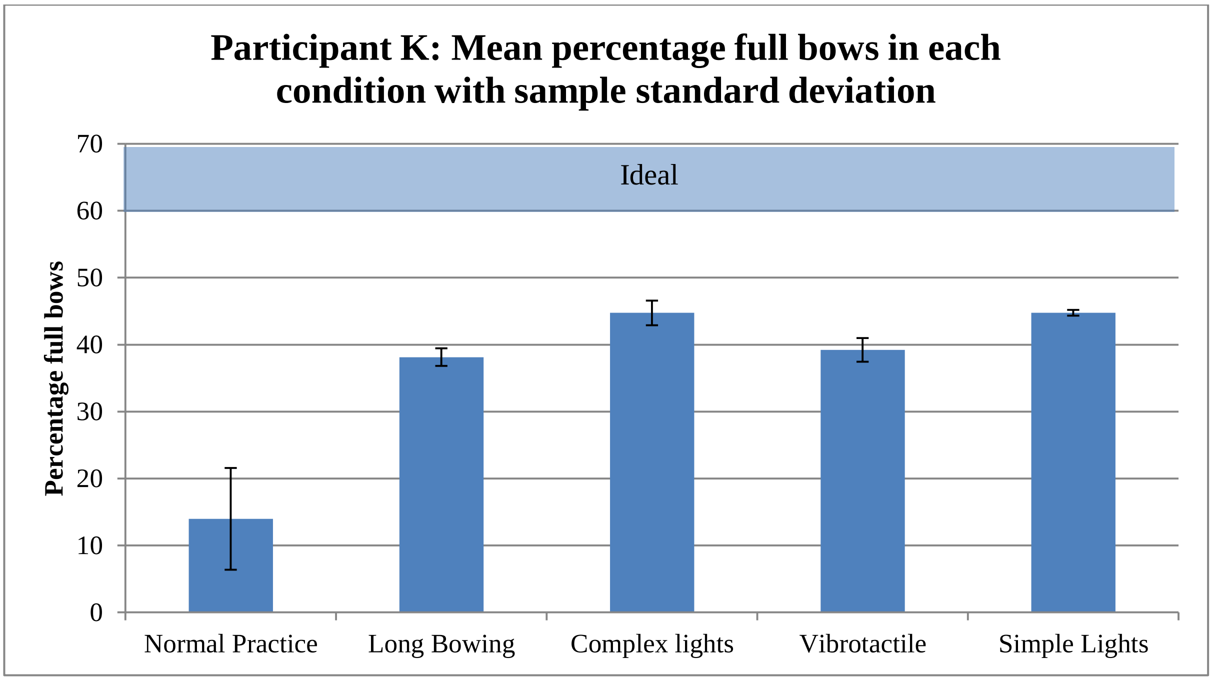

Figure 5-20: Mean percentage full bows used by participant K in each condition. The error bars show the sample standard deviation. There is no after measurement due to a fault in the recording

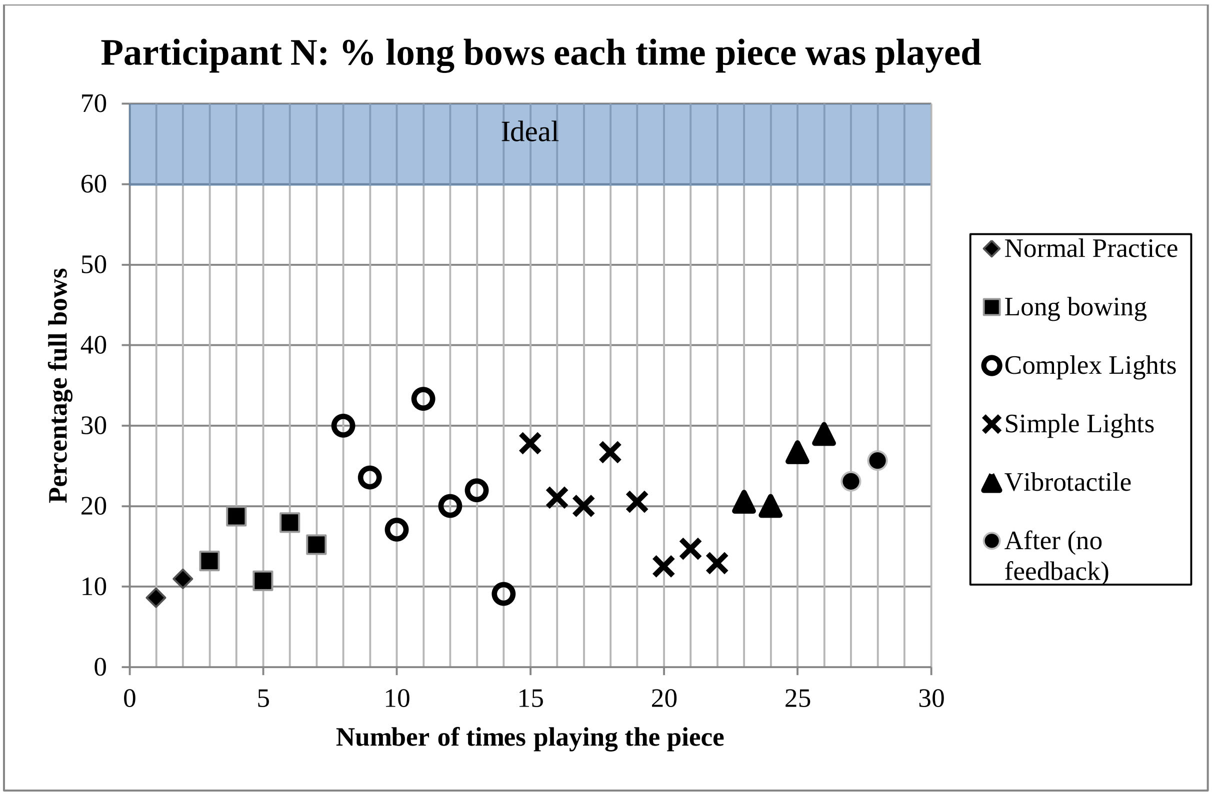

Figure 5-21: Percentage of full bows used by participant N each time she practised the piece. Measurements are taken from the second half of the piece because this is the part she focused her practice on and often began playing half way through the piece.

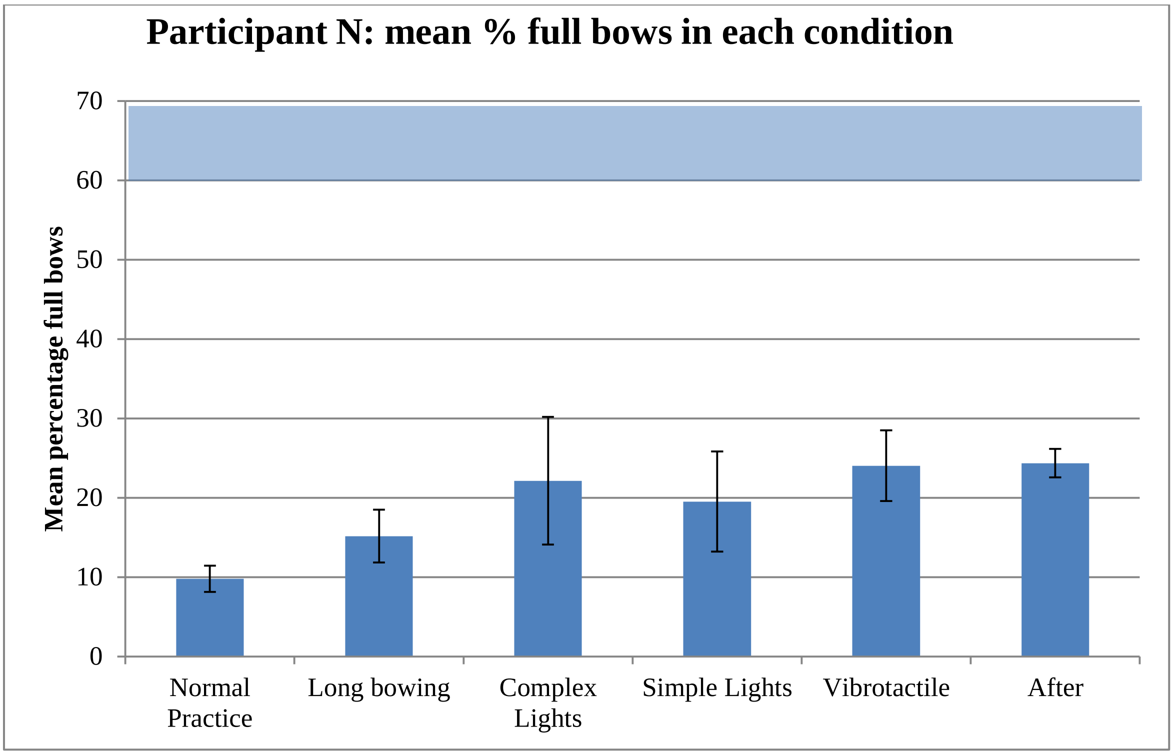

Figure 5-22: Percentage full bows played by participant N averaged over each condition. The error bars show the sample standard deviation.

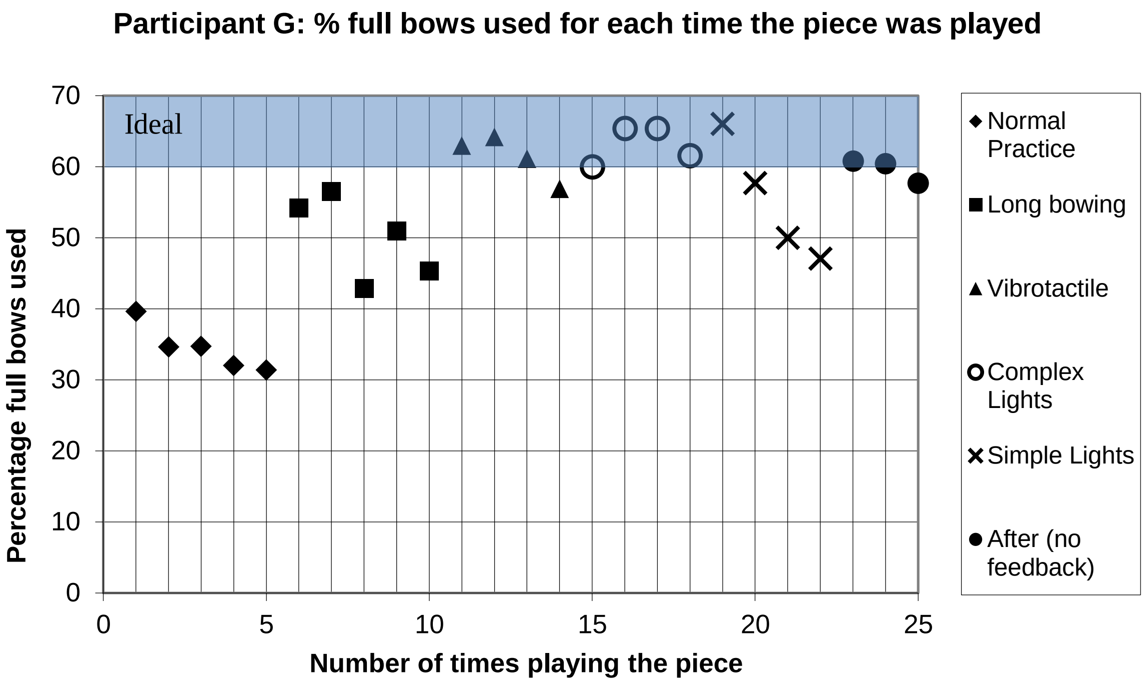

Figure 5-23: This graph shows the percentage full bows participant G used each time he played the piece. As he is using a high percentage of long bows, he may be affected by a ceiling effect which would make it difficult for him to improve beyond the level he plays with the feedback. Notice the falling pattern in the points from the simple lights, this is the feedback he liked least, saying "the flashing" was "too much".

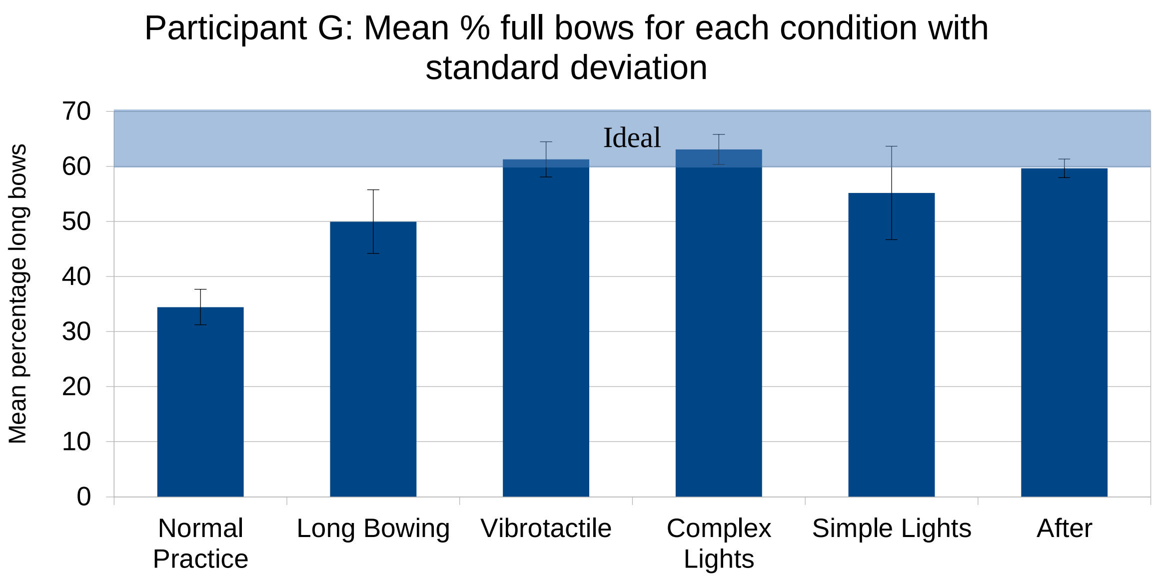

Figure 5-24: Mean percentage of full bows used by participant G in each condition. This is calculated by finding the mean over the repeated times he played the piece. The error bars show the standard deviation.

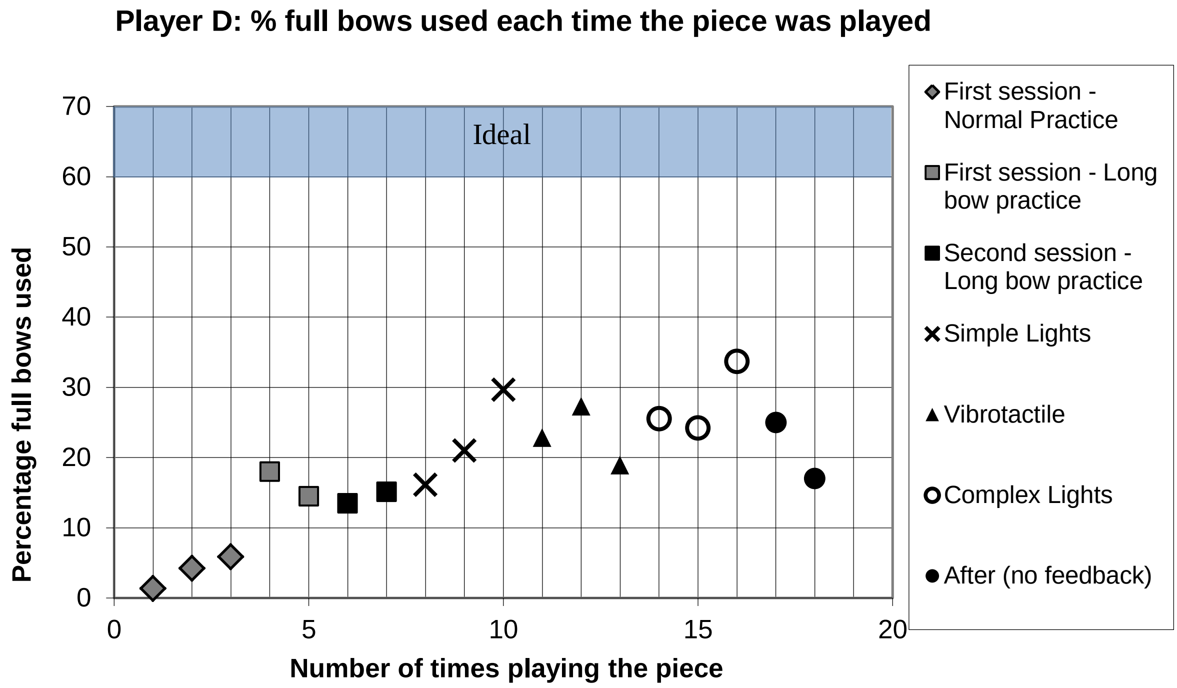

Figure 5-25: This graph shows the percentage full bows participant D used each time she played the piece. This data was taken over two sessions due to the prototype running out of battery in the first session. The long bowing practice was repeated in the second session in order to give a baseline.

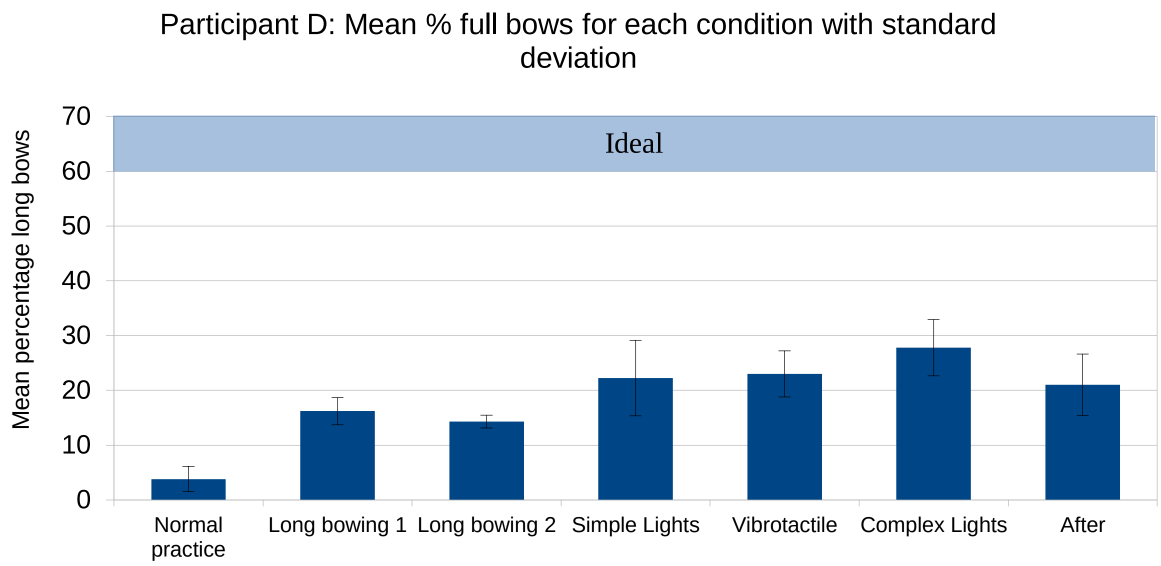

Figure 5-26: This shows the mean percentage full bows played by participant D for each condition. The error bars show the standard deviation.

(iii) General improvement

Participants B and M show a different pattern to the others described so far. Both of their bow usage data shows no preference between the different types of feedback – they were all equally effective. Participant B's data (Figure 5-27 and Figure 5-28) shows that she improved most between her normal practice and the practice condition focused on long bowing. The introduction of feedback then produced additional improvement in the number of full bows she used when playing. It is difficult to say for certain that this improvement is entirely due to the feedback since the data suggests that she was making progress during her long bowing practice so some of her subsequent progress may also be due to a training effect. Although the bow usage data shows no preference for any particular type of feedback, in her interview she did express clear preferences. In particular, she disliked the simple lights because the *"flashing"* was too *"sporadic"* to making it difficult to interpret and would have "*done [her] head in*" if she had used it over a longer period of time. These feelings about the simple lights are very similar to those expressed by participant G. However, in this case they do not seem to have had a direct effect on her performance. She preferred the complex lights because it gave her a *"time-line"* which she found *"easier to use… to determine".* She placed the vibrotactile feedback in the middle because she found it easier to interpret than the simple lights but not as easy as the complex lights.

Participant M shows no improvement between his normal bowing and bowing with the aim of long bowing. The introduction of feedback does appear to have encouraged more full bows but the effect is small. His bowing data (Figure 5-29 and Figure 5-30) shows no difference between the different types of feedback but the interview does show that he experienced them differently. He was most positive about the vibrations because he felt he did not have to concentrate on it and *"it just told me immediately, you know, where I was".* He found the complex lights difficult to use because he had to get his eyes in *"two different positions"* whereas he felt the simple lights could be more easily monitored out of the corner of his eye.

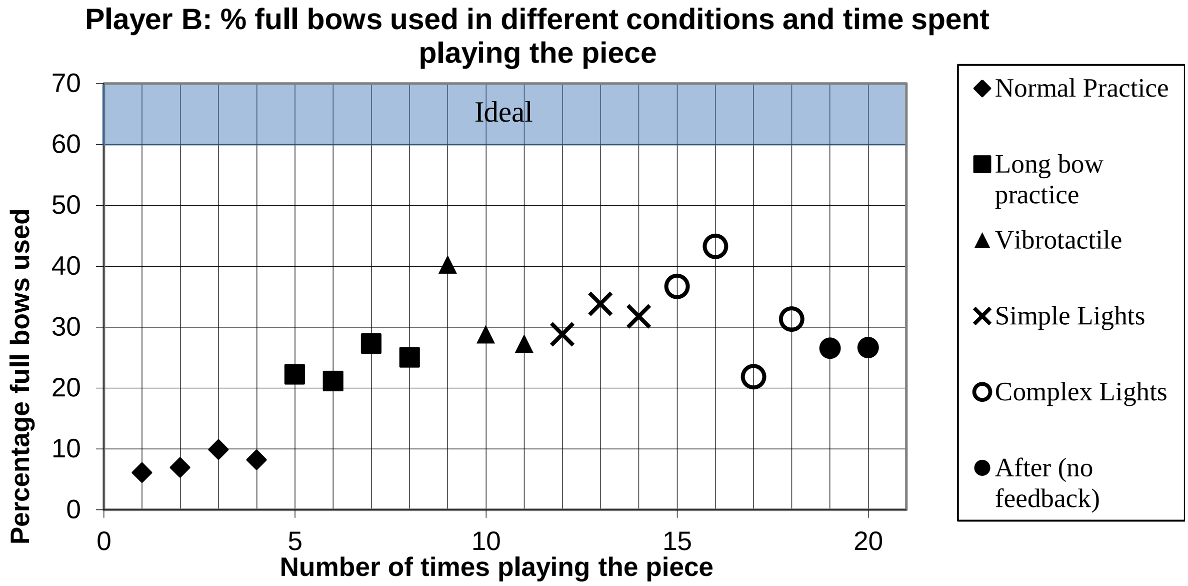

Figure 5-27: This shows the percentage of bows played by participant B which were full bows. Each data point represents playing the piece once.

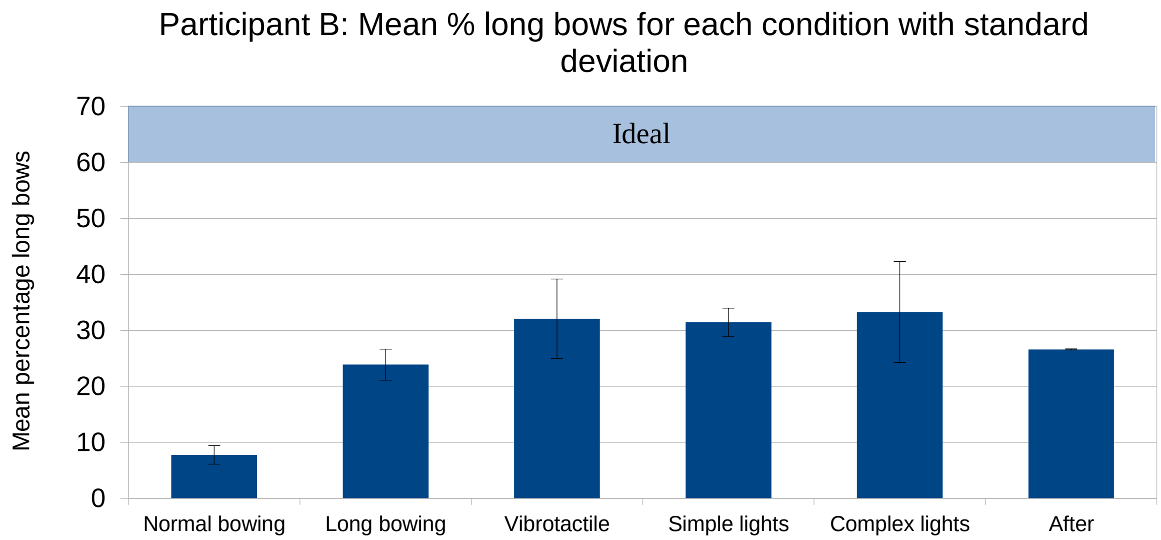

Figure 5-28: This graph shows the mean percentage long bows used in each condition, the error bars show the sample standard deviation.

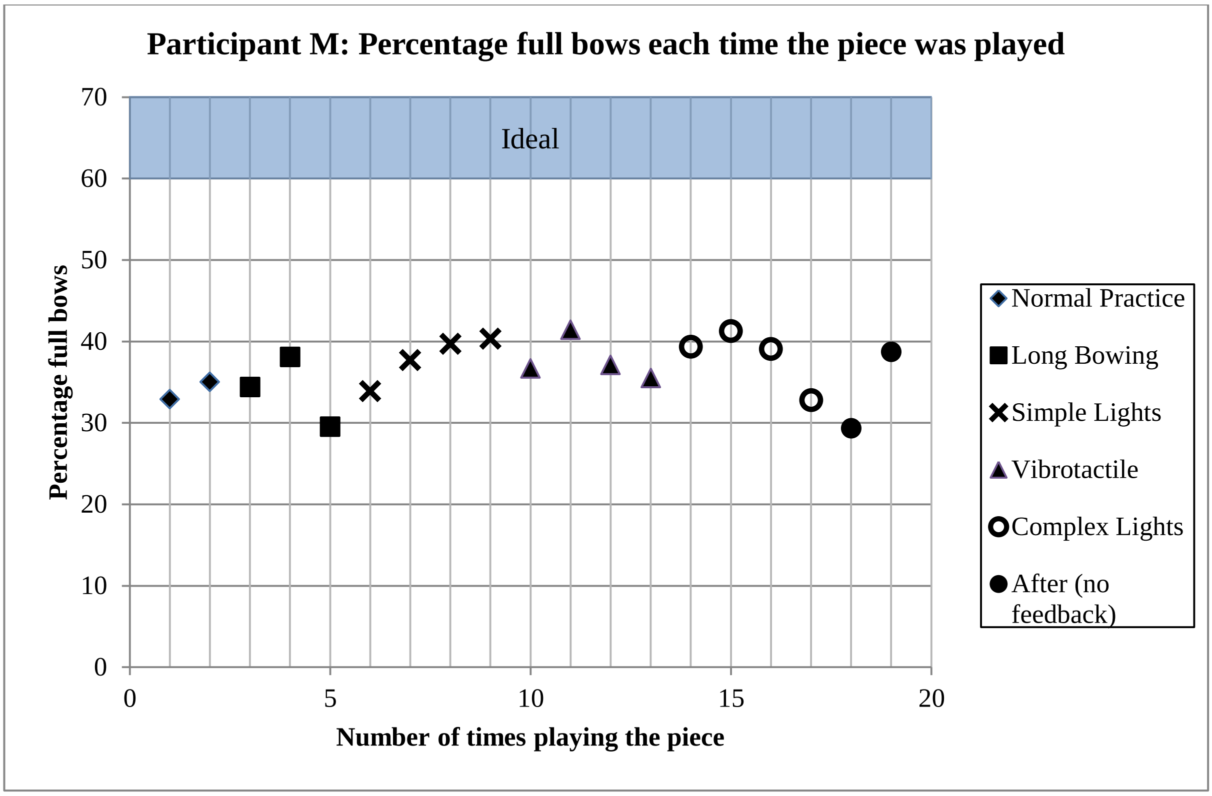

Figure 5-29: Percentage full bows used by participant M each time he played the piece

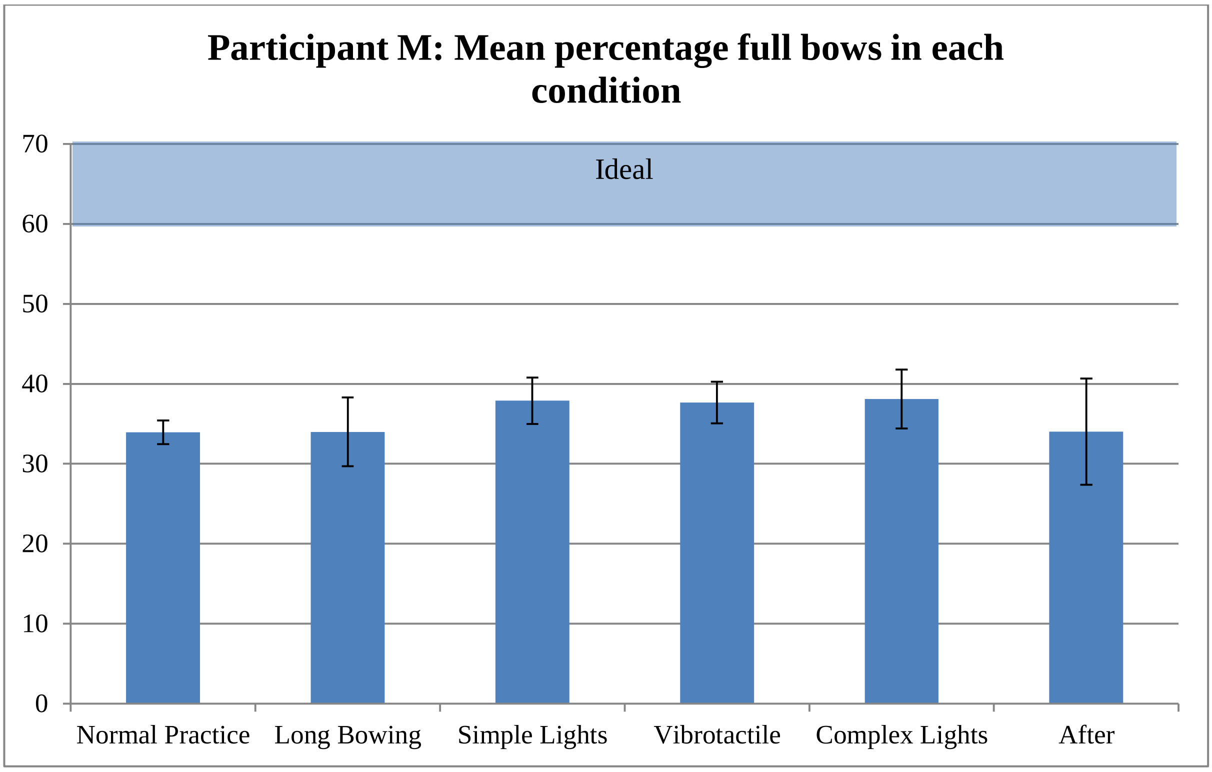

Figure 5-30: Mean percentage full bows used by participant M averaged over each condition. The error bars show the sample standard deviation.

(iv) Changing practice strategy

Some participants responded to the feedback by changing their practice strategy. The most striking example of this is participant E (Figure 5-31). As she started to use the feedback she began to focus in on the parts of the piece where she found it hardest to play long bows and get positive feedback. She is a professional player and thus is keenly able to judge where she feels she ought to be playing long bows and wants to improve where she is not achieving this. In the two feedback conditions with the lights she broke down the piece in this way, only playing a few bars at a time, until she was satisfied that she was playing it with the correct amount of bow as well as intonation and rhythm. The bow usage data on the graph shows that in these conditions she was playing fewer full bows, this was because she had deliberately made the task harder for herself by only playing the parts that challenged her. In contrast, the vibrotactile condition does not show this pattern, this is because she played the piece all the way through in this condition. In the interview she said she disliked the vibrotactile feedback because she found it *"confusing"* and difficult to link to her movement. This may explain why she reverted back to her standard approach of playing the piece all the way through rather than engaging with and breaking up the task like she did with visual feedback. Of the two types of visual feedback, she chose the simple lights as her favourite because she found them *"simpler"*. However, she said that the complex light might be useful to the children she teaches because *"it would give them an incentive to how close to the green they've got"*. This suggests how at different levels of ability musicians may require different things from the feedback. Participant E - who is an experienced player – did not need the extra information that the complex feedback has to offer because she achieved long bows easily when focused upon this task. However, novices may need the extra support of knowing how close they are getting to the positive feedback.

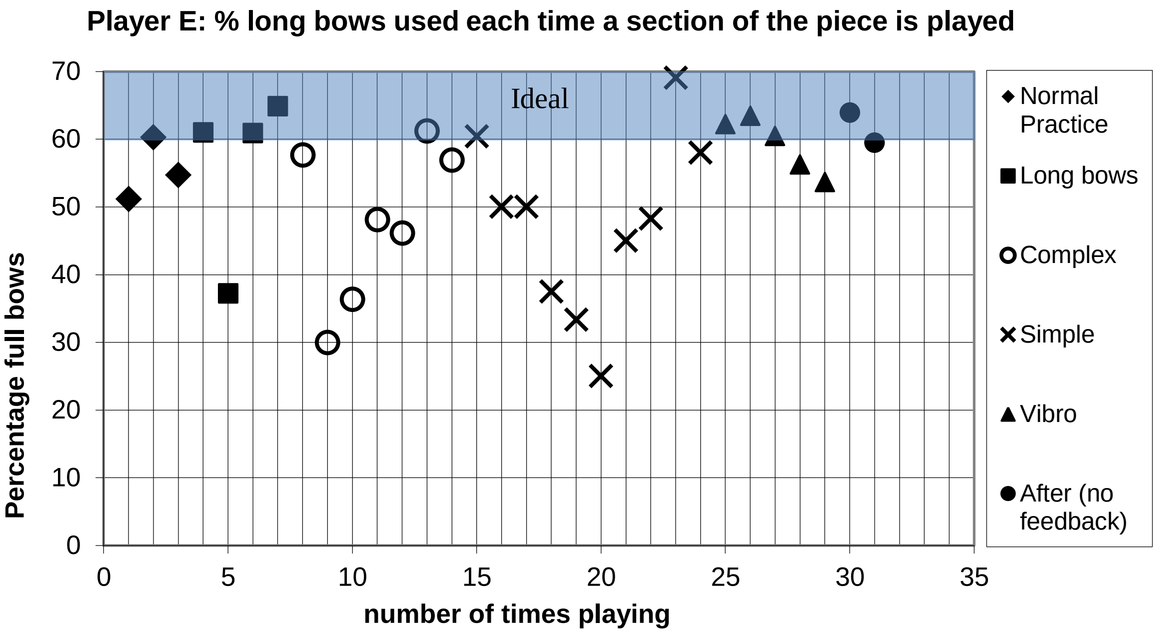

Figure 5-31: Percentage of full bows played by participant E each time she played a section of the piece. Note the points of this graphs should not directly be compared to one another as they come from playing different parts of the piece at different speeds meaning each point represents a different task with a different level of difficulty. For this reason participant E's data has not been combined into a graph of means and was not used when looking for a main effect.

(v) No Improvement

Some participants showed a ceiling effect, which meant that it was difficult for them to improve further because they were already playing close to the maximum number of full bows appropriate for the piece. Participant H, whose bow usage is graphed in Figure 5-32 and Figure 5-33, shows this kind of effect, which prevented the feedback from being effective. After being asked to maximise his bow usage he showed a marked improvement and a high number of full bows for the piece; from then on he showed no further improvement. Participant H is a particularly experienced player and it is not surprising that he should be already able to play the piece with full bows when concentrating on this task. It should be noted that other players (E - Figure 5-31 and G – Figure 5-23 and Figure 5-24) also show a similar ceiling of a maximum number of full bows of between 60-70%. The variation in the exact figure from person to person comes from differences in their interpretation of the piece (exactly which notes they felt needed full bows) and differences in body shape and arm length which slightly changes the exact length of bowing the sensor will classify as a long bow.

Although the feedback had little effect on his bowing, participant H's comments from the interview are interesting. In particular he chose the complex lights as his favourite type of feedback because they fit best with the way he thinks about using the bow:

"to think of length of bow I have to visualise a line – for me that one is good and the blue and green is the most happening line going on there. Literally the line is there and then you go a bit further so I think for me that's the best probably."

He rated the simple lights and vibrotactile feedback equally. He said he found the vibrotactile feedback difficult to connect to his bowing:

"I didn't always feel I was registering what it meant when it was coming. I would feel a buzz and then I'd think oh where am I? Am I half way down am I at the end?"

Although initially saying he found the vibrotactile feedback less visually distracting he also felt it would be best used when he didn't have to concentrate on reading the music perhaps because of the cognitive distraction coming from his difficulty connecting the vibration to the bowing:

"yeah I think for the vibrations to work well for me I would need to know the music – the notes very well. And then I would find it quite handy I think also the physicality of the buzz going on in your arm is good for lightening as well – lightening the weight of the arm – being aware of you know – reminder of the weight of it."

This quote also shows that although he found the vibration difficult to link to his bowing movement he did find the physicality of it useful for increasing awareness of other aspects of his body. Although a very experienced player, he was very self-conscious near the start of the session. This is why there is only a single measurement in the normal practice condition: he started the piece several times, but stopped after a few notes, after this he only had time to play the piece through once before it was time move on to the next condition. He explicitly said on the audio recording that he takes time to relax when playing solo in front of people: "I get very uptight about these things and I need to relax as if there's no-one there really."

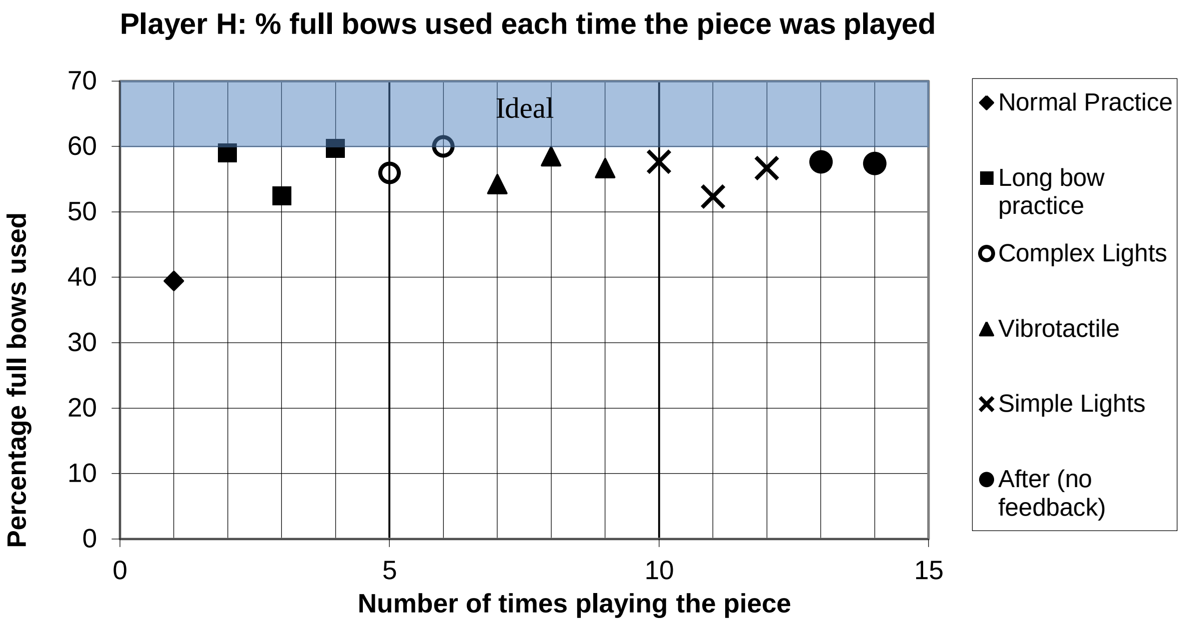

Figure 5-32: This shows the number of full bows used by participant H each time he played the piece. The graph shows a marked improvement when he is asked to use more bow, but no further improvement with the feedback.

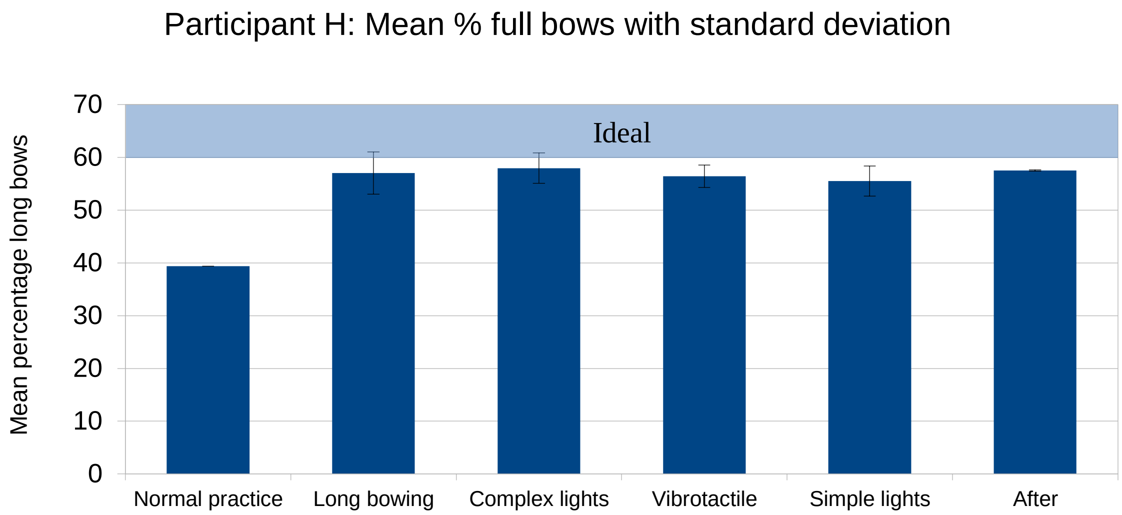

Figure 5-33: This shows the mean percentage full bows played by participant H under each condition. It is likely that this graph is showing a ceiling effect where it would be very difficult for this participant to increase his bow usage further.

Additional interview data

In addition to the thirteen participants presented here with their bowing data, two further participants experienced the feedback but their bow usage data is unavailable due to technical faults with collecting the data. Their interviews were successfully recorded and can be used to complement the data already described.

Participant J preferred the complex lights because *"there was some obvious correlation with where the bow was, and how much you had left to use"* which shows that he found this type of feedback easier to interpret in terms of his bowing. He also described how he felt it was valuable because he had learned something from it:

"I thought I was using all my bow and it was obvious I wasn't even using half the bow at that point so I learnt something from that"

This shows a different type of learning from that envisioned in the design - which was physical type of learning - training muscles through repeating the correct action. The learning he shows here is more analytical, giving him a new idea for an area for him to work on. In particular, he felt that it was easier for him to *"process"* visual signals than the vibrotactile ones and found the vibrotactile particularly difficult to link to his bowing:

"with the vibrations I was too busy concentrating on the notes, I wasn't necessarily correlating when the vibrations started with what direction my bow was going in or how much I was using."

Like participants H and K, this shows a difficulty noticing the vibrotactile feedback while reading music when compared with the visual feedback. This finding refutes our initial hypothesis, which suggested that delivering feedback through a different modality would be easier to use.

Participant O was the first participant to use the feedback. The settings used for his session were stricter than those used by the other participants – the definition of a 'full bow' where you would start getting positive feedback was set to be much closer to the tip or heel of the bow. In particular, the vibrotactile feedback, due to the inbuilt lag time in the motors (the time it takes for them to start moving), came on only at the very tip of the bow which made it very difficult for him to get the positive feedback on this setting. This was not a problem for the other participants. For this reason his bowing data were not analysed. However, the data from the interview is still illuminating although it must be considered in this context.

He chose the complex lights as his favourite because he *"quite liked actually trying to get the difference between the blue and the green – just getting those last few lights to go on"*. This shows how the additional detail in the complex lights which breaks the task down into smaller goals added interest for this participant. He placed the vibrotactile at the bottom of his list, partly because it was difficult for him to trigger it but also he felt it was difficult for him to feel when he did trigger it because it was too similar to the vibration from the lower strings on his viola. This interference issue was also mentioned by one of the cellist (participant F) although he felt that he was able to learn to distinguish between the two quite quickly.

Looking at the data from the participants at an individual level has shown that individual characteristics greatly affect how participants respond to the feedback. The feedback was most successful in cases where participants were unable to improve the length of bow they used through focusing on it alone (long bowing condition). Even in cases where participants could improve their bowing by focusing on it, the support of the feedback appears to enable most participants to improve further. For some participants who started off playing with a lot of long bows it is harder for them to improve, making the feedback less effective for them. For other participants, the feedback seems to encourage them to try different practice strategies and engage with the task in a way they had not before. In the following discussion we use the results from these participants to build a framework about how multimodal feedback acts to change behaviour.

5.6 Discussion

To structure the discussion of these results, a framework is proposed through which to consider the process of using the feedback during playing. This is based on the findings from the study and the theories of perception and action discussed in the literature review. In order for feedback to be effective, interaction with it has to go through a series of stages (see Figure 5.34): (i) it must catch the attention of the player at the appropriate time; (ii) the player needs to be able to take in all the relevant information necessary for the feedback to be meaningful; (iii) the player needs to be able to interpret what this information means in terms of their own playing, and (iv) the player needs to convert this understanding into an appropriate movement.

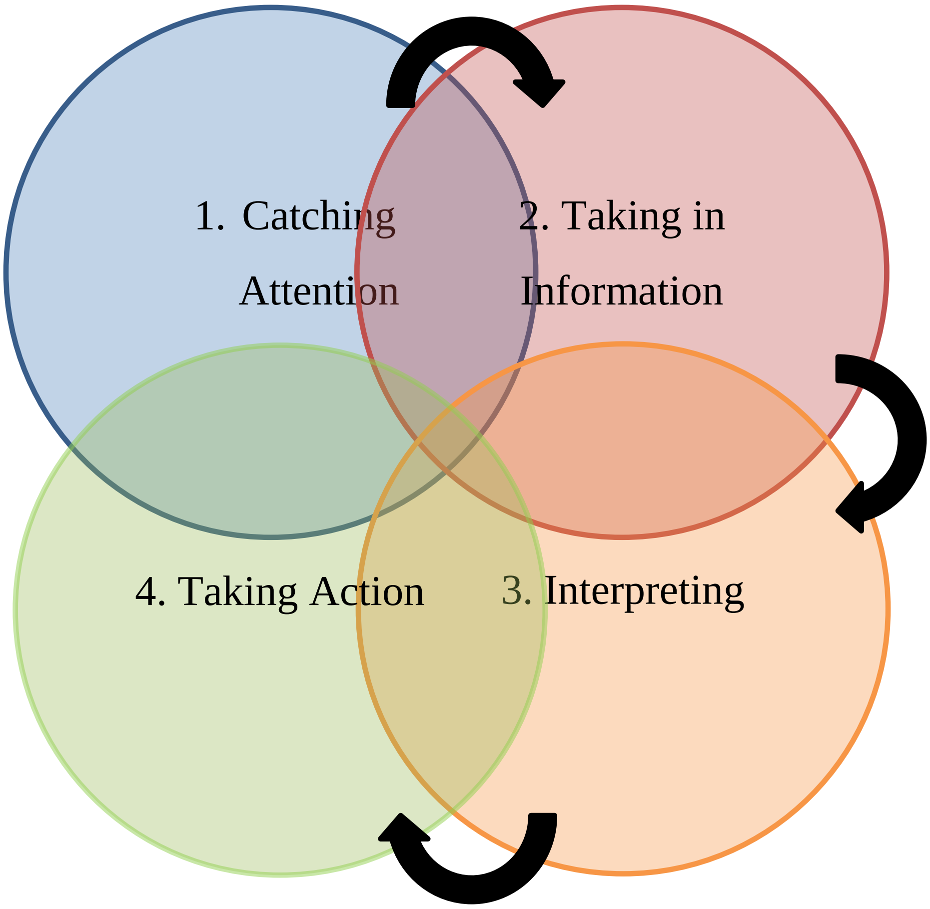

Figure 5-34: Interlinking stages required to use real-time feedback, each stage influencing the others

When a particular feedback was praised or criticised by the participants in the interviews it was often for the way it performed in one of these stages. It is important to note that these processes are sequential, each stage relies on the previous one, for example in order for feedback to be understood, it first has to be noticed and all relevant information taken in by the player. Each stage requires different things from the feedback and, as will become clear as we consider each one in more detail, designing effective real-time feedback involves making trade-offs between the requirements of each stage. It is also clear that while there are commonalities across participants in the way they notice, take in, interpret and act on the feedback, there are also many individual differences as well.

The framework may appear very similar to information processing models such as the three stages of processing used in Multiple Resource Theory (Wickens, 1980): perception, cognition and then reaction. From this point of view, the proposed framework is adding an additional stage of catching the player’s attention - the need for which arises because using the feedback and playing music is an act of multitasking in which the feedback cannot remain at the centre of attention at all times. It is important to point out that the data does not show that these stages are completely separate and there appears to be an overlap between taking in the information, interpreting it and acting on it. Moreover, it seems that the more these overlap the more the participants perceive the feedback. This overlap means that interpretation of our data is also compatible with the Theory of Event Coding (Hommel, Müsseler, Aschersleben, & Prinz, 2001) which suggests that perception and action are closely linked and that we perceive our environment in terms of the actions we are able to perform within it. From this point of view, our analysis of the perception and interpretation of feedback is looking at how we encourage this link between the perception of the feedback and the actions which both created the feedback and can be made in response to the feedback.

Some of the stages in this model are similar to those which Norman describes for overcoming the gulf of execution and the gulf of evaluation (Norman, 1986). However, there are some key differences between the models which come from the differences in the problems they are being applied to. First, in the case of real-time feedback, the technology must play a role in initiating the interaction; this involves catching the user’s attention. This is not a stage included in Norman’s model because his model mainly focuses on traditional interactive technologies where the user initiates the interaction. The stages of taking in and interpreting information are then analogous to the process of negotiating the gulf of evaluation. The stage of taking action is somewhat analogous to the gulf of execution. However, there are key differences here too, with real-time feedback the gulf of evaluation has to be negotiated before the gulf of execution, which is the opposite to Norman’s model. The process of taking action is also different, the user takes action in response to the feedback in order to improve their playing rather than communicate with the technology. In sum, Norman’s model is mainly concerned with how people communicate their intentions to technology; whereas, real-time feedback is about how technology communicates information to people.

5.6.1 The multi-modal real-time feedback framework

Stage 1: Catching Attention

The ability to grab someone's attention is important because the feedback is designed to remain in the periphery, but also needs to be able to become more salient at appropriate times in order to have an effect. Like Weiser's (1991) ubiquitous computers, ideally the feedback would become "invisible" to the player until they need to change their playing and then they would become aware of what they need to do to improve rather than the display itself. However, this is challenging because unlike many of the applications Weiser describes, with real-time feedback the technology must initiate the interaction rather than the user because the player's central attention is heavily focused upon playing. Thus, the feedback must bring itself to the player's attention to become useful. This is a step further than being "ready for use at a glance" which would be the ideal in cases where the timing of the interaction is less crucial or there is less in the user's central attention.

One thing that is clear from the findings was that the two types of visual feedback differed in their ability to attract participants' attention. The simple lights - because they flash on or off to communicate a single discrete variable (full bow or not) - were considered to be better at attracting attention than the complex lights which constantly move. This suggests that sudden changes in a display are more attention grabbing than the more incremental changes.

Some participants found that the complex lights were unable to attract their attention and for them it made them unusable, whereas in contrast they found the simple light much easier. On the other hand, there were participants who found the complex lights caught their attention to an appropriate level, whereas the simple visual lights attracted their attention too strongly, distracting them from their playing. This shows that while there is a general trend for the simple visual feedback to be more noticeable than the complex visual, how this affects a particular individual is different. These individual differences may be dependent upon many things, such as a participant's music reading abilities or how aware they are of things in their peripheral vision. The second of these is particularly interesting because it is something orchestra players need to be able to do, as they watch the conductor and read their music at the same time.

One of the ideas behind our original hypothesis - that vibrotactile feedback would be more effective than visual feedback - was based on the idea that perceptually it would be a more appropriate channel for catching attention and communicating information. This is because it is a different modality to reading music. In terms of gaining a participant's attention, it was true for only some of the participants. Many who rated the vibrotactile feedback highly did indeed comment that not being visual made it less intrusive and distracting. However, the sensor results showed that in some cases it became so unobtrusive that it ceased to be effective. In contrast, one participant found it caught her attention too strongly and another found the sensation uncomfortable. This shows that there are marked individual differences between participants in the way vibration catches their attention and suggests that perceptually some people may be better suited to using this modality than others. This has also been suggested by other research, for example Nagel et al. (2005) and Pielot and Boll (2010) .

Stage 2: Taking in information

Having noticed the feedback, participants then need to be able to take in all the necessary information to make the feedback usable. This consists of not only seeing or feeling the display but also being aware of the position of their arm and in the case of the discrete signals knowing when the signal went off.

The complex lights display the most information as part of the feedback. These were criticised by some participants because they found there was too much visual information to take in while also reading the music. On the other hand, other participants found the complex lights made a stronger link to their own movement which meant that although the display is complex, they had to put less effort into simultaneously being aware of their body.

The vibrotactile feedback and the simple lights both give a single signal to the player and as such they do not require the player to take in much information directly from the device. However, in order to use the feedback the players must also know when or where the signal came on while they were playing. The initial hypothesis suggested that vibrotactile would have an advantage because it is positioned on the arm and this should help to make participants more aware of that it is the movement of that limb which is producing the feedback response. This was true for some individuals who said they found it easier than the visual signal to link to their movement. One participant favoured the simple lights over the vibrotactile because she found them easier to use to make links with the written music. Rather than linking the feedback directly to her movement, as we had envisaged, she shows a preference to link it to the written notes which her movements are a response to.

With both discrete displays (e.g. vibrotactile and simple lights), participants reported difficulties determining with the necessary accuracy when the feedback came on. This perceptual problem made this particular type of feedback difficult for them to use. Different participants reported this difficulty with the simple lights compared with those who had difficulties with the vibrotactile. It seems participants may have difficulty linking discrete displays to their movements. Hence, simply selecting one modality over another will not solve this for all individuals.

By comparing the complex and discrete displays, it can be seen that a balance has to be found between giving too much information through the feedback so that it is difficult to take in and giving too little so that it relies too heavily on the player to take in information from other sources to make it meaningful.

Stage 3: Interpreting

Interpretation and perception are closely linked. In order to parse the relevant information, the player needs to have some understanding of what the feedback means and then decide what is relevant. Equally important, in order to interpret the feedback, the player needs to have perceived it in a necessary amount of detail.

Considering the complex lights, no participants complained that they were difficult to interpret in terms of bowing. The main problems participants had with them were from not being able to notice and take in the display. Several participants initially liked using the complex lights until they came to the task of playing from written music. Some participants gave additional details about why they found the complex lights particularly easy to interpret: one participant talked about it fitting with the way he visualised the bow as a line and another likened the display to time-line. This suggests that this visualisation was tapping into pre-existing knowledge that the participants had, which made it easier for these participants to use the feedback. The ease which participants found in interpreting this display is likely to be because it continuously gave feedback in response to the player's movements, in terms of Norman's gulf of evaluation (Norman, 1986), the time between user action and feedback from the system is minimised so that users are able to make a clearer causal link between their movement and the outcome of the feedback. Moreover the mapping between the length of the line and the length of the bow is simple and direct, with a single input mapping to a single output with both being in the same dimension, length.

One of the initial ideas when hypothesising about which types of feedback would be more effective was that the complex lights would be more motivating because they show the incremental steps towards a goal. This idea was supported by three participants who commented specifically about aiming to get a step further or liking to see how close they could get to the green zone of the lights. The lights were talked about in terms of incentivising by another participant, who talked about seeing the green light as 'good'. This participant did not see the vibrotactile feedback as praise in the same way and no other participants described the feedback as an incentive or a reward. This suggests that visual feedback may be more suitable for the expression of praise and goal achievement than vibrotactile feedback.

Stage 4: Acting

Once the feedback has been understood, it then has to be transformed into an appropriate action. Our initial idea suggested that vibrotactile would be best for a longer bowing action to players because it is positioned on the arm where this action is made. One participant did specifically say this was the case for her, although her bow usage data showed longer bows with the simple lights feedback so it suggests that this impression may not be accurate. More generally we would expect feedback that makes a clearer link between the bowing movement and the responding feedback would be better at encouraging further long bowing action. There are a variety of individual differences in how this link is made which have already been discussed in the interpretation section.

Another way to examine this is to consider the mechanism through which the correct action might be promoted by the feedback. This needs to be considered carefully because apart from the complex light display participants do not get a response from the system unless they have reached their goal and made the correct movement, rather than the correct movement being a response to the feedback.

One way in which the feedback may be working is that participants are directly responding to the feedback signals in real-time. Participants may do this by continuing to move their bow, only changing direction when they have felt or seen the feedback to tell them they have gone far enough. If this is the case, the timing data might show that when participants initially use the feedback the playing tempo would be slower compared to without as they are spending more time over each note. This was observed in some cases, most notably participant A. Another way participants may adapt their movement in real-time is that if participants have not triggered the feedback in the way they wished over the last few bows they may speed up their bow movement in order to fit in larger bowing movement without slowing the tempo.

A third way the feedback may change playing action is to enable participants to pick out areas of their playing where they are not using much bow and then to work on these. This can be seen in the data from participants E and F and in the interview responses from J and K who both talk about how the feedback allowed them to notice notes or phrases where they were not bowing as they intended. Interestingly, three of the four participants who show this behaviour most clearly preferred visual feedback and the fourth participant, although preferring vibrotactile feedback, identified the section he wanted to work on while using the simple lights.

A fourth way feedback may change bowing action is that it incentivises long bowing and makes it a higher priority when playing than it was previously. This was seen in the discussion of the data from participant F as well as being suggested by comments from other participants such as participant O.

A fifth way to change their playing action is that it changes their focus of attention from an internal one to an external one. As discussed in the literature review an external focus, which looks at how movement affects something outside the body, produces better performance than an internal focus which dwells purely on the movement of the body itself (Weiss, Reber, & Owen, 2008). For example it may be that participant A, who was unable to improve the length of his bowing through focusing on this as a task was focused too internally on his own movements when bowing, and that externalising the focus was the way in which the feedback acted to create such dramatic improvements in his bowing. It is hard to say which feedback would be best for supporting an external focus, participant A used vibrotactile feedback first which transformed his movement but he also played well with visual feedback as well. Therefore it may be that any modality of feedback can act as an external focus in the right situation.

5.6.2 Summary

The study conducted here did not statistically support the two hypotheses and suggests that no modality is universally superior when it comes to giving real-time feedback. Instead, there are many individual differences in the way people respond to different modalities and designs of feedback. There were also some indications from the individual data that different modalities might be better suited to different mechanisms for encouraging good bowing. For example, visual feedback may be more suited to helping participants pick out areas where they need to improve, whereas vibrotactile feedback may be better for use for real-time adjustments to movement. This supports the idea that participants should have access to different modalities of feedback in order to support and encourage different practice strategies as well as catering for individual preferences.

This study examined real-time feedback in individual practice. The next study investigates how the use of real-time feedback changes when it is used by an ensemble playing together. This raises new questions about the differences between visual and vibrotactile feedback because visual feedback can be shared whereas vibrotactile feedback cannot.

5.7 A Final Word on Methodology

This study used a mixture of data sources, including sensor and interview data. When conducting a study like this the rapport that is built between researcher and participant is important. As I knew the participants beforehand it made me reflect upon my role in the study; how my relationship with the participants affected the running of the study and how building a relationship with participants through participating in joint activities might be advantageous. Building on existing personal relationships with participants during recruitment setup an atmosphere of familiarity, which continued into the study and affected the way participants interacted with me. For example, criticism of the devices was often phrased in a constructive way, giving me enough detail to enable improvements or balanced with a positive comment. This kind of interaction shows participants are responding to me, both on a personal level and as a data gatherer and that they recognise that I had an emotional interest in the long-term success of the design.

Knowing the participants and playing with the orchestra also helped me to recognise the challenges they faced in the study. For example nearly all participants seemed nervous and many said they felt like one of the weakest players in their section of the orchestra. I could empathise with this nervousness; playing on your own is revealing something very different about yourself to playing in an orchestra. I offered reassurance by admitting that I couldn't play these parts myself and joking that maybe the conductor had chosen the hardest parts to get us to practise. My presence is likely to have contributed to this nervousness. However, participants may have felt equally nervous about playing for a recording which they knew would be listened to and analysed. At least by being present I could try to alleviate the nervousness by offering reassurance. Being a member of the orchestra gave me the facilities to identify with the emotions of the participants and offer reassurance in a way a stranger could not. It also facilitated a more informed discussion of the music as I was familiar with the repertoire.

These meta-level reflections also raised the question of authority. The learning focus in this study needed to be relevant and useful to participants outside the study itself in order to make practice meaningful to participants. To choose such tasks requires authority that comes from a high level of musical knowledge. However, I did not have this authority as I was musically equal to the participants. For this reason it was very useful that the tasks themselves were sections of the current repertoire selected by the conductor as being relevant. This meant I could use the conductor's authority to justify the need for long bows without the need to claim any access to special knowledge myself. This is something I used during the study as many participants pointed out that this was not how they would choose to play those sections otherwise.

Chapter 6 – ShareSense: Studying ensemble playing through participation

6.1 Introduction

This chapter describes an in-the-wild study of a system developed for helping adult musicians play together. The setting was a summer school where they come together for an intensive week of practice and playing together. A system was designed to provide both (i) personal feedback in the form of real-time vibrotactile feedback and (ii) group feedback in the form of a dynamic visualisation that each contributed to and could see. The work builds on the research findings of the previous study by investigating the visual and tactile modalities for providing feedback. However, instead of comparing one against the other it explores how they could be used together for different purposes: vibrotactile for individual feedback to aid playing and visual feedback to help players keep abreast of how they are doing when playing together. This line of research extends from an additional observation from the previous study, which was that many of the orchestra players commented that they felt they were the weakest player in their section of the orchestra and it made them feel stressed. This may or may not be true; there isn't an objective way for them knowing if this was the case. This observation led to thinking about whether a real-time shared visualisation could be designed to provide shared information that would let each member see their contribution to the overall ensemble. Although in some cases this could make a less able participant feel worse; there is also potential for feedback to encourage players to support one another, rather than only focusing on their own short-comings. The goal of this study, therefore, was to explore the second research question from section 1.3:

> *How can both shared and personal real-time feedback support playing as part of an ensemble?*

The previous study contrasted vibrotactile and visual feedback to aid individuals practicing long bowing. The findings showed that both vibrotactile and visual feedback can help people improve their playing but they are perceived differently by different individuals. In particular, some individuals appreciated the instinctive understanding that they felt the vibrotactile feedback gave them, whereas others preferred the more analytical information that they could deduce from the visual feedback. The goal is now to see how it can be used by a group of players simultaneously. A shared display could potentially encourage them to discuss their playing with one another. Referring to the framework from chapter 5, shared feedback may help with catching attention because players could point out things that other players have missed.

However, the additional complexity of a shared display may make it difficult to take in all the information while playing. Having the individual vibrotactile feedback may help to combat this as it only gives information about the individual which should be easier to take in. In terms of interpretation, shared feedback gives a different baseline for comparison because players can compare themselves with one another. This may motivate players to take action in order to make their contribution to the ensemble more equal or having public feedback in this way may make players feel more uncomfortable about their role in the ensemble.

6.2 Research Aims

The research aims of the study were to explore how giving feedback to players may influence the experience of playing in an ensemble and encourage mutual support. Real-time feedback provides the opportunity to represent every person's contribution to ensemble playing in a shared way and potentially helps people to recognise their own contribution and comment on others.

Designing feedback for ensemble playing, however, presents new design challenges compared with providing for individual use. Ensemble playing involves additional coordination skills; namely being able to synchronise their movements with each other that places additional visual and cognitive load on the players. It requires different members of the group taking the lead at different points (Glowinski, Coletta, Volpe, Camurri, Chiorri, & Schenone, 2010). If feedback is going to help them play, both in terms of helping to correct their individual playing and coordinated ensemble playing then it needs to be not very demanding. The aim was to design an ambient visual display which gives information in real-time to show participants their contribution towards a shared goal relating to bowing. For ecological validity it is important that the members of the ensemble use the feedback towards a shared learning goal that is relevant to their playing as an ensemble and that they are motived to improve.

6.3 Methodology

An in-the-wild study was employed to study the different kinds of real-time feedback for aiding ensemble playing. The setting that was chosen was a summer school for adult amateur musicians where attendees have actively sought out the opportunity to play music with others. It presented an opportunity to study feedback in an intensive way; there is a range of students from different musical backgrounds who are all there for week with the intention of playing music every day. However, this setting has its drawbacks, running the study at the summer school meant relinquishing control over many details of its design as the organisation of the school could not be changed for the study and so it had to be adapted to meet the needs of the school and the participants.

To understand the experience of playing in an ensemble with real-time feedback a participant observation method was employed. Participant observation is widely used in social sciences and anthropology, most commonly associated with ethnography (Hammersley & Atkinson, 2007). It provides researchers with an insider perspective allowing them to interpret their observations with an intimate understanding of the values, beliefs, social norms and experiences of the group or culture being studied. In the previous study, I had benefited from being a member of the same orchestra as the participants. This helped me to build rapport with participants and empathise with the challenges of the music because I had experienced them, too. The methodology in this study takes this idea of participation and shared experience a step further, participating in the same activities as the participants both during study sessions and outside of them.

The reason for this choice of method was to gain an insider perspective to help me interpret my observations and conversations with participants. It can be strongly argued that we interpret the world and the views of others by drawing on our own experience (Rode, 2011). By experiencing what it is like to use the prototypes playing with the participants I can gain useful experience to interpret the comments that they give me. By also participating in the setting of the study (a musical summer school) I learn about the culture and values of the participants, which will also aid with interpreting data. Gaining this insider perspective and shared experience will also aid building a good rapport with participants.

6.3.1 Study Setting

The study took place at a music summer school attended by approximately 200 adult amateur musicians. The week long school comprised rehearsals for the first five days and a day of concerts on the sixth. People could take part in a variety of courses, including classical orchestral playing, folk and jazz. Participants for the study were recruited by an email sent to all the string players who had subscribed to the summer school. Six players volunteered to participate, four violinists, one viola player and one cello player. All were unknown to me. I had also never attended the summer school before. Once recruited, participants were asked by email about their musical interests and the courses they were attending. The majority of participants were interested in classical playing, but one of the violinists was a jazz player and another played exclusively folk music. Therefore, it was important that the prototype was versatile enough for groups playing a variety of stringed instruments and playing a variety of musical styles.

Five of the participants had attended the summer school before. One, however, had not and it was hard for her to judge whether she would have time to dedicate to the study. The participants had courses scheduled for nearly all the hours between 8am and 8pm with only one or two hours break a day when they might be able to meet for the study sessions. Sessions were scheduled with participants in advance for approximately 40 minutes every day. The participants met with me either as a pair or singly depending on timetable constraints. During the sessions it was planned that they would play duets or trios and I would join in to complete the ensemble.

Participants were asked if they would like to suggest music to play in the study sessions. The folk player suggested some repertoire for her session and pieces of music were exchanged by email in preparation. The other participants did not suggest pieces so I brought a selection of duets and trios. Repertoire from the summer school courses was also used.

6.4 Design of ShareSense

A feedback prototype, ShareSense, was designed for the study based on the following criteria:

- **Flexible Use** – It should deal with a range of instruments found in an ensemble (such as violins, violas and cellos and a variety of playing styles)
- **Adaptable** – It should be possible for players to adapt it to set different goals for different pieces
- **Group Feedback** – The feedback must have a group element which shows the ensemble as a single entity
- **Individual Feedback** – Individual players should be able to identify their contribution to the group and use the feedback to actively pursue their personal goals for improvement
- **Quick to Set-up** – Participants should be able to start using it within minutes
- **Quick to Learn** – It should not take long for participants to learn how to use

6.4.1 Sensing

To determine which aspects of playing to give feedback about, the prototype needed to be *flexible* enough to be used on multiple types of instrument and with many styles of music. Bowing is an action common to violinists, viola players and cellists. It also plays an important role within ensemble playing. It is usual where possible to try and synchronise bowing of players in an ensemble and to try and use similar amounts of bow as one another. It was decided therefore to further develop the prototypes for long bowing for use with the orchestra players. The previous study showed that length of bow could be simply and effectively sensed in all these cases using a single gyroscope. An armband was built for each player similar to those used in the MuSense prototype. These sensed the length of the players bow stroke and sent this wirelessly to a laptop which controlled the visual feedback. The vibrotactile feedback was built into the same armband as the sensor.

6.4.2 Feedback

Real-time feedback was given in two modalities: visual and vibrotactile. The visual feedback provided information about an individual player relative the group as a whole. The vibrotactile was designed as an individual private form of feedback that enabled players to monitor their own personal goals. These two sets of feedback could be used together to give a multimodal experience.

Visual Feedback

A new visualisation was designed that showed all the players bow strokes in one display. This visualisation was projected onto a large screen in front of the players (see Figure 6-1, Figure 6-2 and Figure 6-3). The visualisation was designed for two or three people playing at once because these were the types of group that were available for the study at the musical summer school. Each bow stroke was represented as a petal-like shape which expands over time and becomes longer with the length of the bow used. The intention is that each player produces petals of a particular colour on a particular axis when bowing. Bows in one direction form petals on one side of the flower; bows in the opposite direction form bows on the opposite side of the flower. The flower-like layout was designed to be simple to be read at a glance, enabling players see if they were using similar amounts of bow to others by looking for symmetry. If the petals were asymmetrical it meant they were playing at different lengths. The flower design was also intended to give the sense that everyone's bowing when combined made the overall shape – displaying individual petals as part of a group effort.

The design draws on ideas from gestalt theory (see Wagemans et al., 2012, for a review) where the whole is different to the sum of its parts, meaning it is possible to grasp something as a whole object without being consciously aware of the details of its constituent parts. In this case, the symmetry and shape of the flower as a whole can be understood at a glance. Whether a group of shapes is perceived as a whole object is dependent upon the functional relations between the constituent parts. Symmetry and synchrony are two examples of these functional relations. Synchrony in the movement of petals will occur when participants are moving their bows at the same speed. The symmetry of the flower shows how similar each individual's bow use is relative to the others in the group. Bow strokes not symmetrical to the rest of the group will stand out as being separate to the whole.

When someone plays their next bow stroke, the petals from the previous strokes do not immediately disappear. Petals from the last four bow strokes remain on the shared screen, with fading in colours, to allow people to practise bowing patterns. The presentation of information in the visualisation was intended to be non-judgemental and facilitate different interpretations by the group depending on their aims. Some examples of using it for different possible aims are outlined below. For example:

- **Synchronisation:** In some pieces, the ensemble need to all keep their bows synchronised by using a similar amount of bow. The visualisation was designed to enable the group can track the synchronicity of their playing by looking at the symmetry of the flowers they are drawing (Figure 6-2a).
- **Long bowing:** In some pieces long bows are required to give a smoother, louder, richer sound. Players can compete with one another to play with consistently long bows and can easily compare themselves to the others by looking at how long their petal is compared to the overall flower (Figure 6-2b).
- **Short bowing:** Some pieces require short bows to enable the player to play quickly or to articulate notes. Players can compete to play consistently with short bows by looking at how short their petal is compared to the overall flower (Figure 6-2c).
- **Special bowing:** The shape of the petal also represents bow speed (fat petals are slow bows, thin petals are fast bows). Therefore, special bows will make particular shapes, for example bow that is supposed to be loud and then quiet (fast and then slow) will be a narrow at the bottom and wide and rounded at the top. Players can pick out particular shapes and use them to discuss and compare their playing.

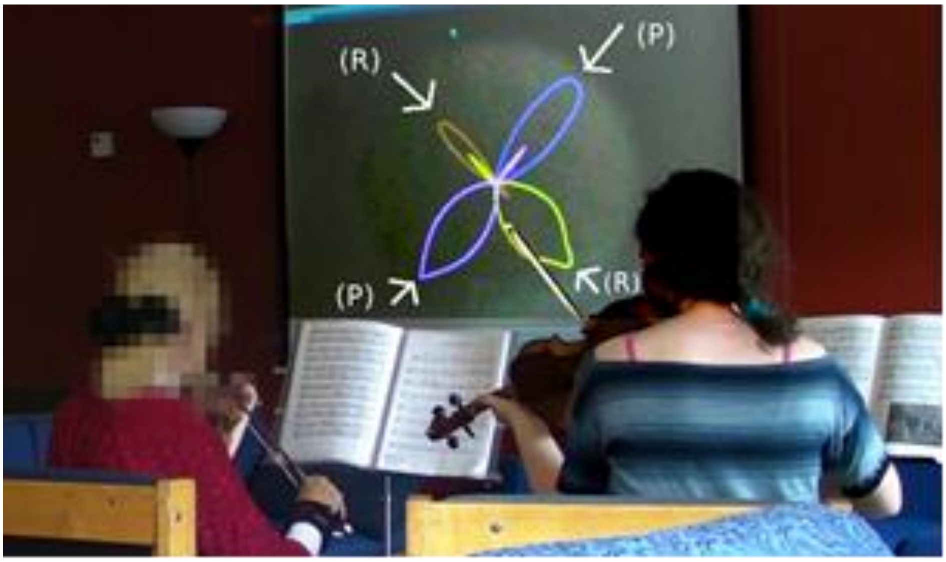

Figure 6-1: Participant (left) and Researcher (right) playing together using the visualisation. (R) = petals made by Researcher, (P) = petals made by Participant.

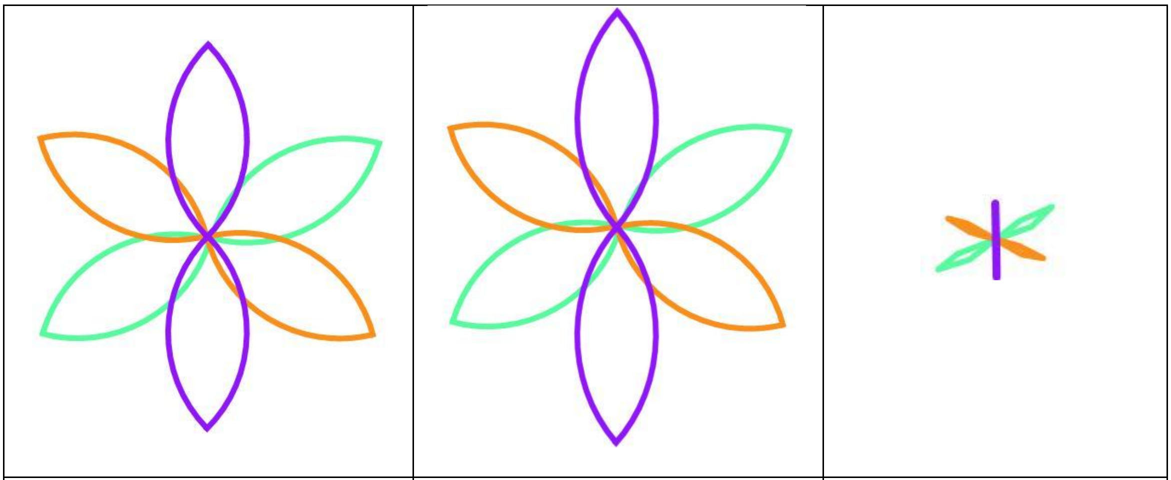

a) Synchronisation – if all players are playing at the same speed with the same length bow then the flower will look symmetric.	b) Long bowing – Players can compete to use the longest amount of bow. In this picture the purple player is using longer bows than the others.	c) Short bowing – Players can compete to play short quick bows. In this picture the blue player is playing with the shortest bows.

Figure 6-2: Examples of how the visualisation can be used for different aims

Vibrotactile feedback

The vibrotactile feedback was not made into a shared display because the vibrations were seen as personal, being felt on the individual's body, therefore making them unsuitable for a shared display. The vibrotactile display was the same as that used in the last study except for one difference: participants could set the length along the bow that the vibration would be triggered. To this end, the armbands were fitted with a knob which could be turned to make the adjustment. This feature was added so that participants could set the vibrotactile feedback to help them meet their own bowing goals. The participants could decide whether to interpret the vibrations as a goal to be reached or as something to be avoided. For example, when playing a quick piece with lots of short bows they could set the vibration to come on after a short length and aim to avoid triggering any vibrations. This more *adaptable* design meant that it could be used for any piece that was to be played. So instead of feedback to encourage long bows, the vibrotactile feedback was designed so that the individual could calibrate their own bow length goal. This form of player calibration also takes into account the differences in arm length between participants. Adjusting this knob changed the scale of the player’s petals on the visual display to keep the two types of feedback consistent with one another and to make the visual feedback conform with the new sensing calibration. So long as all ensemble players calibrated their sensor band with a shared aim then the symmetry of the flower would still reflect the symmetry of their movement.

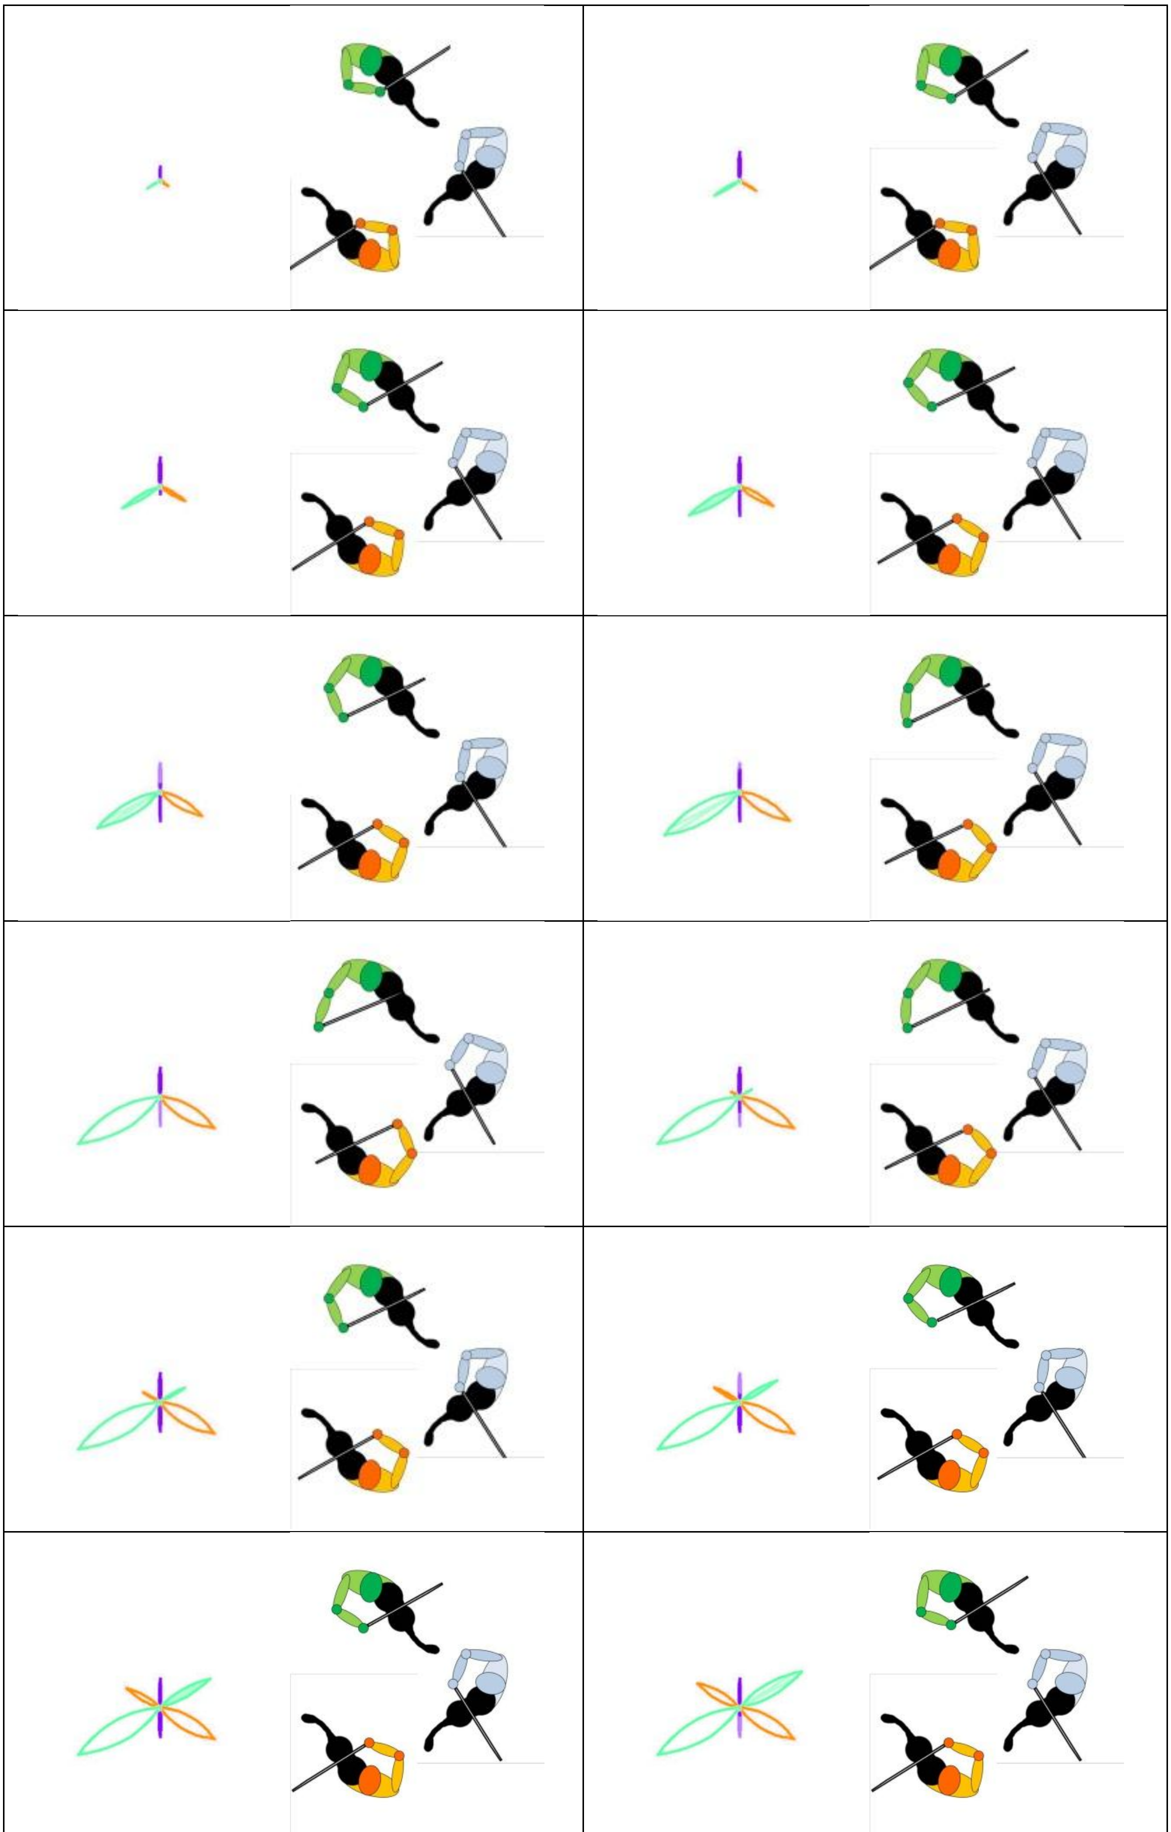

Figure 6-3: An example of a trio playing with the feedback. The green and the orange players are doing synchronised long bowing; the green player is using more bow than the orange player hence his petals or longer than the orange player's. The purple player is playing short, quick bows hence the short, narrow petal shapes.

6.4.3 Implementation

The movement of the bow arm was sensed using a single axis gyroscope positioned on the top of the lower arm. The gyroscope was linked to an Arduino Pro Mini (Arduino, 2014) which took the readings from the gyroscope and integrated and smoothed them to infer the length of movement used in each bow stroke, as previously described in chapter 6. Controlled by the Arduino was a 10mm vibration coin motor of the same type used in the previous study (Precision Microdrives, 2014) which was activated when the inferred length of movement passed a certain threshold. This threshold was set by a variable resistor which was also attached to the Arduino and whose resistance could be set by turning a knob. This was put into a potential divider circuit with one of the analogue inputs on the Arduino being used to read the potential difference across the resistor. By turning the knob the user changed this reading. The Arduino used the reading from the variable resistor to set the length of bow that needs to be played before the vibration is triggered. Therefore turning the knob one way resulted in very short bows and turning it the other way resulted in longer bows.

The Arduino was also linked to an XBee series 1 radio transceiver. Via this it sent the current length of bow used to another XBee transceiver attached to a laptop. Before sending, each measurement was adjusted based on the setting of the variable resistor so that each measurement was scaled to the length of movement required to trigger a vibration. When multiple armbands were in use they took turns to send their data and identify themselves using a character such as ‘a’, ‘b’ or ‘c’. The Arduino in band ‘a’ sets up this turn taking when it is first switched on. Initially, it was planned that the transceiver connected to the laptop should control the data flow by requesting the reading from each band individually in sequence. However, this was found to be too slow, creating a lag in the update of the visualisation. This is because the visualisation had been written using Java which was ill-suited to serial communication with the XBee transceiver. So instead the Arduino in the sensor bands controlled the turn taking by sending a trigger for the next band after they had finished sending their data.

The laptop received the bowing data from all the armbands available and drew the visualisation by plotting the bow length in polar co-ordinates to form the petals of the flower. The current bow length was transformed into the length of the petal. The rest of the shape of the petal was made from the readings of bow length from previous points in time plotted symmetrically either side of a centre line of the petal with the angle from the centre line being equivalent to time.

6.5 Study

As mentioned earlier, the study took place during the summer school. It ran on the summer school site with the permission of the organisers but was not part of the official summer school courses. Instead, participants came to the study between the courses they were attending. The goal was to investigate how the participants used and took to ShareSense in an ensemble, and to find out whether shared real-time feedback can be both useful and usable for ensemble playing. A vibrotactile prototype was built for each participant to use. The visualisation was set up for each practice session.

6.5.1 Participants

Six adult amateur musicians agreed to take part in the study: a jazz violinist (Andy), a folk violinist (Cathy), two classical violinists (Danielle and Eleanor), a classical viola player (Fiona) and a classical cellist (Beth). Pseudonyms have been used to protect anonymity. Participants were recruited by an email sent out by the organisers of the summer school to all string players who had signed up to the school. Seven participants volunteered but one had to withdraw before the study due to ill health. Once the participants had volunteered, the researcher kept in email correspondence with them individually up to the start of the study, using this as a means to learn more about the participants and to give more information about the study. On the first day of the study, participants had the study explained to them again and signed a consent form (appendix F.3) to say that they agreed to take part and could withdraw at any point if they chose.

All participants were experienced players: some having taken up their musical instrument later in life (Cathy, Fiona and Beth); some returning to playing after a long absence (Andy had learnt a school then stopped playing for 20 years before becoming a jazz violinist). Some played violin as a second instrument (Danielle's first instrument was flute). All participants played for pleasure and some also liked to perform (Andy in a jazz band, Cathy in a folk band and Eleanor as a soloist in one of the summer school concerts). All were attending the summer school and taking part in several courses so they had very little free time. Unfortunately Beth found that participating in the study was too much in addition to all the courses and chose to drop out after the first session.

Participant	No. sessions attended	Played together with…	Instrument	Preferred musical genre
Andy	4	Fiona for one session	Violin	Jazz
Beth	1		Cello	Classical
Cathy	5		Violin	Folk
Danielle	4	Eleanor for two sessions	Violin	Classical
Eleanor	2	Danielle	Violin	Classical
Fiona	2	Andy for one session	Viola	Classical
Myself (researcher)	15	Everyone	Violin	Any

Table 6-1: Summary of the participants who took part in the study

6.5.2 Procedure

Participants met with me in ones and twos for a number of 40 minute ensemble playing sessions. The aim was to meet every day for each of the five days of the summer school but this was not always possible due to tiredness and other commitments (Table 6-1 shows how many sessions each participant attended). Participants came to the sessions either individually or in pairs and I played with them as well to form either a duet or trio respectively. Although these are only small ensembles, they are still very different to individual playing because participants have to hold separate parts against one another and can be heard by the other person/people they are playing with. The small groups were advantageous from a study point of view because they allowed the researcher to hear the opinions of all participants. The sessions were held in the common room of one of the accommodation buildings. In these sessions the feedback was introduced progressively. In the sessions on the first day only the vibrations were used. From the second day onwards the visualisation was used as well. This was to allow for a contrast to be made between giving feedback to the individual and giving feedback to the group. Whenever participants used the feedback I did as well in order to share the experience of playing in the ensemble with the feedback. When feedback was first introduced, participants had time to experiment with it to ensure that they understood what it did. At the start of each session participants were encouraged to adjust the calibration of ShareSense so that the vibrotactile feedback was triggered in suitable place for their learning goal. This also meant that they checked that the vibrotactile feedback was functioning and discernible before starting to play.

The repertoire we played varied depending on the interests of each participant. With Andy we played jazz standards which he improvised on and later folk tunes (irish jigs). With Cathy we played folk duets some of which we arranged before and some from the courses she was attending. With Danielle and Eleanor we played trios and then the parts from the string ensemble course (I played 1st they played 2nd violin). With Fiona we played some jazz with Andy and then some of the parts from the orchestra course (viola and 2nd violin had very similar parts which needed short bows). With Beth we played a duet for violin and cello. As there was no set objective for the feedback participants had to choose how they wanted to use it. All started off with it set on long bow to encourage more bow usage for a richer more confident sound. Later in the week some used it for short bows for fast pieces like the theme from 'Cowboys' (John Williams) and traditional Irish jigs.

6.5.3 Researcher Role

In this study I took the role of both researcher and participant. As a researcher, I explained what the prototypes did and possible ways to use them. I also organised all the sessions and provided music as possible repertoire to play if participants did not volunteer their own. During the sessions, to gather data I asked the participants how they found using the prototypes and reflected aloud on my own experience using them.

As a participant I joined with the other participants to play music during the study sessions and at courses at the summer school and informally in the evenings. I used the prototypes during the study sessions and experienced what it was like to try and be conscious of real-time feedback while playing in an ensemble. I experienced the nervousness and pressures involved in playing in a small group as well as the sense of shared achievement when it went well. Through conversations with participants both in the study and more generally I learnt about their approach to music and the pressures they were under during the study, such as a heavy schedule, back pain and a concert they were playing that evening. The role of being a participant gave me new insights into the experiences of what it meant to take part in such an intensive and high-pressured group event. However, it also came at a cost; constant participation throughout the week was exhausting and being in the thick of the action made it sometimes hard to stop and reflect more objectively on the progress of the study.

6.6 Data collection and analysis

Three methods of data collection were used: an observation diary, videos and a questionnaire. Between sessions and in the evenings I recorded my observations and feelings in a diary which I sent to my supervisors and this diary was extended shortly afterwards with additional comments after reflecting on my experience. All sessions were video recorded to give an external view of the events of sessions which I would not be able to see whilst in the thick of playing. Unfortunately due to constraints of being at the summer school, which made it difficult to set up the cameras, some videos are unclear both visually and in recording the comments and conversations, however they are useful as an aid memoire. At the end of the study participants were also asked to fill out a questionnaire comparing the two types of feedback, five of the participants returned their answers.

The notes made during the study were reviewed at the end of the study and used as a basis for the video analysis. All videos were viewed and further notes were made about what happened in each session and the conversations between playing music. From the notes made at the time and the additional notes from the video, themes were identified. These were mutual support, setting objectives and calibration – and are discussed later. Transcriptions were made from selected parts of the videos to illustrate the themes and to study them in more detail.

6.7 Findings

We begin by addressed the central research question:

> *How can both shared and personal real-time feedback support playing as part of an ensemble?*

Overall, the findings from this study showed that shared feedback can support a sense of being part of an ensemble and encourage players to discuss their playing and support one another (see Mutual Support later in this section). However, it depends very much on contextual factors as to its efficacy. From the study, most of the participants commented on how difficult it was to use the shared visual display while playing. One example is Eleanor's response in her questionnaire:

"While playing music it was not possible to look at the visual feedback. So would only look at it when our group was trying it out and playing some scales with long and short bows. […] It was not 'off-putting' but I could not be 'aware' of it while playing therefore limiting its purpose."

Cathy also commented on the difficulty of using the visual feedback while trying to play written music.

"Because I was almost always playing tunes which I needed to follow the music I could not look at it most of the time. It was interesting to see and not off-putting."

I also experienced this; I found I needed to make a conscious effort to look up at the visualisation rather than it attracting my attention. In some cases, participants tried using the visual feedback in their central visual focus. For example, Danielle and Eleanor appeared to enjoy the visual feedback most when using it for scales where they could focus on their effect on the visualisation.

Hence, the shared visual feedback had difficulty in catching player's attention and making them aware of their own and each other's playing. One explanation for this is the ensemble playing setting which is more demanding than the individual practice studied in the previous chapter. When practicing in an ensemble, players must keep in time with one another. Whereas the participants in the previous study of individual practice would stop and start and replay parts where they had made mistakes, this occurred much less in ensemble playing. While participating in the study, I noticed that the pressure to keep time with the others made playing more demanding. It was particularly demanding in cases where I had to play a part which was very different to the parts of the others around me. Ensemble playing places extra cognitive and perceptual load on players because they have to listen and watch the other players to stay in time with them. As such there is less capacity available for monitoring feedback.

Hence, real-time peripheral feedback, in the form of visual feedback, may not be well suited to this type of ensemble playing. Ensemble playing at this level may be too difficult a task to add extra demands on player's attention. Instead it may be better, as some of the participants themselves suggested, to wait until players are confident in what they are playing and then introduce feedback as the central focus.

How effective was the vibrotactile feedback for personal playing?

Many of the participants also failed to notice the vibrotactile feedback while playing – even more so than in the previous study. This is despite the checking that they could feel it when playing simple exercises. Next we consider why this was the case. Four out of the six participants reported not noticing the vibrotactile feedback and some felt that they almost never felt it while playing. For example, Andy commented in session three:

"I didn't notice it. Particularly while improvising I don't notice it."

Beth put in her questionnaire that she did not notice it while playing.

"I almost never noticed it while I was playing!"

On the other hand, Cathy believed *"I always felt it"*, Eleanor also said she had *"had no difficulty with the vibrotactile feedback."* So as found in the previous study there is a large variation in people's perception of the vibrotactile feedback. However, these are even more pronounced than in the last study because participants are reporting never feeling the vibrations while playing whereas in the study of individual practice, the participants just reported having difficulty noticing it.

The vibrotactile feedback was set to indicate when a goal has been reached. Not feeling the feedback can confuse players about how well they are doing in achieving their objectives. In some cases participants may be led to believe they are not reaching their goal when they are, in others participants may believe they are not feeling the feedback when in fact they are not using enough bow to trigger it. One particularly confusing case happened with Beth, who found it difficult to be aware of the feedback but also did not regularly use enough bow to trigger the feedback. She began by saying that she was not been aware of the feedback and outlined two possible reasons why this might happen:

"To be honest I've been almost totally unaware that I have this until I stop and I move my arm more dramatically and then I feel it, but it's almost like I only feel it when I stop. Now that may be because it's been set at longer bow strokes than I've been playing."

This comment demonstrates the difficulty for participants of trying to interpret not being aware of feeling any vibrations. It could either mean that she was not using enough bow length or it could mean that she was not aware of the vibration because she was so heavily focused on the music. Following this she tried playing some scales in order to feel where it came on. First she played a scale with short bows and then longer bows and commented:

“Now that’s interesting because that I could tell when I wasn’t using as much bow. I think partly I was thinking more about other things [when playing the previous piece]”

This shows her concluding that she probably just didn’t notice the vibrations. However, the video showed that perhaps the reason she did not feel the vibrations was she wasn’t using enough bow to trigger them. When she started to play the scales she also started to increase the amount of bow she was using. Therefore, the vibrotactile feedback was unsuccessful in this case of making her more aware of her bowing use. Eleanor and Cathy were able to monitor the vibrotactile feedback while playing but could rarely look up from the music to see the visual feedback. Conversely, for Andy the visual was more useful and noticeable. He also expressed a lot of difficulty noticing and interpreting the vibrations. The large cognitive load which ensemble playing places on people seems to polarise them more strongly in terms of their ability to notice vibrotactile feedback, making it useful for some people and not others. These individual differences in peripheral awareness are summarised in the last column of Table 6-2.

To reflect upon when and how the real-time feedback can be useful in shared settings, three themes of mutual support, setting objectives and calibration are outlined next. These were chosen by reflecting on what were the differences between groups that controlled the success of a particular type of feedback. Success was judged based upon my observations and feelings as a member of the group about whether a group were playing well together and whether the feedback was playing a useful role in that and on what other members of the group said at the time and in the questionnaires.

	Mutual Support	Setting Objectives	Peripheral Awareness
Andy	Videos show that we played together but did not comment on each other's playing. I used the feedback to talk about how I could improve but Andy did not.	Andy generally set long bowing objectives, but he commented that this was an arbitrary decision made for the sake of the study as he had no difficulty with bowing.	Andy commented that he had difficulty noticing the vibro-tactile feedback, particularly when improvising. The questionnaire shows he felt more aware of the visualisation and liked it because it gave detail on speed as well as length.
Beth	Video shows we supported each other by discussing the piece and the bowing, but did not comment on each other's playing.	Beth worked on long bowing but commented that she was not sure that this was the right objective for her on this piece.	Beth commented that she was not aware of the vibrotactile feedback, but video data shows that she was not playing long enough bows to trigger it.
Cathy	Video and notes show that we supported each other working on long bowing based on what we could see in the visualisation and discussing the vibrotactile feedback. We did make comments about each other's bowing.	In the first session Cathy and I identified that she had difficulty with long bowing and agreed to work on it together over the week using the feedback.	Cathy commented at the time and on the questionnaire that the vibrotactile feedback was a good reminder and helpful to her. She said she could not look at the visualisation while playing but sometimes looked at it at the end of a piece.
Danielle & Eleanor	The video shows a supportive ensemble group with some detailed discussion about the pieces. Visual feedback helped to break the ice and was discussed when playing scales. Mutual support mainly based around note learning not feedback.	Long bowing and short bowing objectives were chosen as a group based on appropriateness to piece. Main focus was on learning the notes of the orchestral music not on bowing. They commented that feedback would be better suited to practice exercises.	Video and questionnaires show that Danielle had difficulty noticing the vibrotactile feedback, while Eleanor had no difficulty being aware of it. Eleanor said she found it difficult to be aware of the visualisation.
Fiona	My notes show a difficulty feeling like a coherent group playing with Andy. Fiona asked if we could work on course pieces in separate sessions. Video shows mutual support doing this but not centred around feedback.	Together we chose to set short bowing objectives for the fast orchestra music we were learning. However, the main focus was on learning the notes of the orchestral music not on bowing.	Fiona reported that in general she was not aware of the vibrotactile feedback.
Me (researcher)	In some groups I felt confident to offer support to others based around the feedback. In others I did not because I was self-conscious of my own failings.	I followed the objectives of the other participants, but I was surprised that I was not meeting the long bowing objectives and wanted to improve this.	I had some peripheral awareness of the vibrotactile feedback but not all the time. I was not able to be aware of the visualisation unless I made special effort.

Table 6-2: Summary of the way each participant used the feedback based on observations from the videos and notes made at the time of the study, comments made by participants in the sessions and the questionnaires filled out afterwards . Danielle and Eleanor are grouped together because they played together in the study and at the summer school.

(i) Mutual Support

By mutual support we mean discussions and advice shared between players to help them improve their playing. Individual differences were found between the groups as to whether the feedback encouraged mutual support. When playing with Cathy, in particular, the visualisation did add to the experience of playing in an ensemble. We were both aiming to use longer bows in our playing. During our sessions playing together we both referred to the visualisation at the end of a piece from time to time, for example:

"*Woohoo look at that lovely big petal!* [Cathy]"

Sometimes, we also talked about the vibrotactile feedback as well saying whether we had triggered it. Occasionally we commented on each other's petals in the visualisation, for example, at one point Cathy pointed out, with some pride, that her petals were bigger than mine. This made me want to work harder on my bowing. The shared visual feedback allowed us to make comments about each other's feedback in this way which we would not have done otherwise. As well as referring to the feedback we also supported each other in other ways, such as discussing and planning the bowing for a piece to play together. In this way we supported one another in encouraging each other to reach for our goals. At the end of the study, Cathy was very positive about her experience. She felt that the feedback had helped her to improve her bowing:

"I found it a very worthwhile experience and feel my playing has improved as a result. Particularly using all the bow, and keeping to time."

One of the main reasons why we worked so productively together is that we both shared the same objective, namely, to increase our bow usage. Therefore, we could use the visualisation to compare each other's progress. We were also playing at a similar level, she had the advantage that she was playing familiar pieces, but I had more experience, making us evenly matched. The visual feedback may have also been more accessible to us as well. Cathy was mainly playing pieces she already knew and so may have had more opportunity to look up from the music than some of the other participants. I also felt less dependent on the sheet music in the sessions with Cathy as we mainly played folk music which lends itself to being memorised and played by ear.

Another case where the feedback added to the sense of being in an ensemble was during the second session when playing with Danielle and Eleanor. In the first session, with only the vibrotactile feedback, we had found it difficult to engage as a group. In the second session, two factors came together to make us feel more like a group. Firstly, we had played together that day in the orchestra course and we used this session as an opportunity to practise this repertoire which meant we had a motivating task. Secondly, this was the first session where we used the visual feedback together. Danielle and Eleanor were able to understand how to use the visual feedback more quickly than they had with the vibrotactile display. We began by playing scales together and watching the petals of the flower move together and then playing scales at different speeds and watching the different patterns it made. This helped us think about playing in time with each other. Once the session progressed to playing the orchestral music we did not discuss the visual feedback while we were playing. Instead, we were much more focused on playing the correct notes with the right timing. The visual feedback was not able to play a role in this and we had no spare attention for anything else. However, during that session we worked well as a group, making useful comments to one another. The exercises with the visual feedback at the start of the session had played a role in breaking the ice and making us feel more connected to one another.

These two examples show that visual feedback can support ensemble players in a different way to vibrotactile feedback. The case of Cathy in particular shows us helping each other to achieve a shared goals through discussion based around what we saw on the display. However, the fact that this did not happen very much with the other groups shows that it depends on a number of other factors. It seems most likely to be of benefit when the members of the ensemble have a shared objective when using the feedback and are all interested in the music they are playing. Being well matched in terms of ability may also make people more comfortable discussing their progress with one another. It is also necessary to be aware of the feedback, which was a problem that came up for both types of feedback.

(ii) Setting Objectives

Setting objectives helps people interpret the feedback. In this study, it became clear that to use ShareSense effectively, an objective had to be set by the ensemble to use the feedback together. The visual feedback was found to work best in cases where there were clear objectives which were well motivated and chosen by the participants. When participants were unsure of what objective to choose or felt that their choice of objective was arbitrary, the feedback became largely irrelevant. The following examples show how setting objectives varied between different individuals and how this affected the use and usefulness:

Andy came to the study with good bowing technique and a particular bowing style suited to his jazz playing. Quite early on he identified that he needed to decide to what his objective was for the piece he was playing in order to make sense of the feedback and he decided to aim for long bows. However, he felt that this objective was somewhat *"arbitrary"* rather than being something he felt the piece needed. The video showed that after this he did play using more bow. However, towards the end of the first session he still felt that *"I don't have a clear objective in mind and that's the trouble"* and he used this to account for his experience of not noticing the vibrotactile feedback *"even though it was working"* and he was using enough bow for it to be triggered.

The findings also showed that for high level players it can be hard to set a clear objective because bowing becomes a matter of style and expression which may not be easily categorised. In later sessions, for example, Andy enjoyed using the visualisation because it showed more of the characteristics of bow use and where the petal shapes changed with bow speed and length. Rather than giving himself a simple objective of long or short bows he explored the visualisation looking at the different shapes he could make through his movement and suggested a different use for the feedback altogether – namely performance art. In his case this makes the feedback more usable in two ways – first it gives a more engaging level of challenge and intricacy and second it brings the feedback into the foreground and makes it part of the performance which overcomes his difficulties in being aware of the feedback while performing.

In the study sessions, Danielle and Eleanor, and Fiona used ShareSense to practise the orchestral music that they were learning for one of the courses. They were able to choose bowing objectives that were appropriate for the pieces of music. However, their main aim in these practice sessions was to learn the notes of the orchestra music. Therefore, ShareSense did not give feedback on their main objective. This made it hard for them to use it and meant that the feedback was often overlooked once the sessions were underway. They felt that the feedback would be better suited to practice exercises where they would be able to give more attention to bowing.

In contrast, Cathy started the study with a wish to improve her bow usage which she felt was a weakness for her. It was something she had identified as a problem in the emails she sent me before coming to the summer school. Her objective in the sessions was to use more bow. Over the sessions the videos showed her using fuller bows with the feedback. She also noticed an improvement and felt that the feedback was helping her to achieve this as already discussed in the previous section.

Beth was also prone to using quite short bows. She came to realise this on starting to use the vibrotactile feedback when she used the feedback to analyse her bowing:

"So I obviously only use the middle of my bow because I didn't feel it at all. So that goes to show something… But I don't know if this piece would have long bow strokes."

The second part of this quote shows that her response to realising that she was not using much was not to aim to lengthen her stroke but to question whether she needed to. In this piece, although some of the notes may have only needed short bows there were certainly parts where full bows would have been much more appropriate and the video reveals that she was only using half the bow. However, she did not make it her objective to try and lengthen her bow. This shows that although feedback can tell a person something about their bowing, it is down to individual judgement as to how to use this information to set objectives.

(iii) Calibration

The third theme concerns the effect of enabling the participants to adjust when to receive feedback. For many, it became apparent that a knob was not a good mapping for making this adjustment. Firstly, it was not clear which way to turn the knob to lengthen or shorten the bow length. Secondly, there was no quick feedback to see what change they had just made. Participants had to turn the knob, then play a bow to see where the feedback came on, then turn the knob again and test again until they found the setting they wanted. Perhaps as a consequence of this, participants did not take much ownership of the calibration process and only adjusted the calibration when changing their objective from long bowing to short bowing when moving on to particular pieces.

6.8 Discussion

The findings from the study showed that ensemble playing in a real world context is very demanding and that providing both personal and shared feedback to help groups play together and be aware of how they are doing relative to each other's playing is limited. Below we discuss these in relation to when and how it might be used to good effect.

6.8.1 Can real-time feedback be designed to support the sense of being part of an ensemble?

The sessions with Cathy showed that a shared visualisation can help people to support one another in their learning by facilitating people to be aware of the other player's bowing as well as their own. However, this kind of mutual support around the feedback was infrequent. This suggests that the focus of the feedback had to be relevant to all members of the ensemble and the ensemble needs to choose a group objective that they can support one another with. Second, participants needed to be familiar enough with the music that they could dedicate some attention to working on this objective. Third, players need to have an interest in the music they are playing. Therefore while it is possible to support people playing together it is a challenging design aim and the context of use needs to be carefully chosen.

It is also important to notice that the mutual support shown in the sessions was towards shared personal learning goals. The feedback helped us to learn about and improve on our personal playing technique and support each other in doing this; rather than helping us to improve our playing together as an ensemble. This may be due to the situation – we knew we were not going to perform together – therefore our goal was to improve our playing for when we would play in other ensembles.

The findings from this study show that there are advantages in providing shared feedback. The sessions with Cathy showed us using the shared visual display to support one another by making comments about our bowing. However, it is also important to note that we also made comments about the vibrotactile feedback to one another and this was not a shared display. One difference though is that in the case with the shared display we can comment on each other's playing and this can potentially helped us to use the feedback more effectively and may also be motivating.

The session with Danielle and Eleanor also shows how a shared visualisation can break the ice and encourage people to play together. The case of Andy using the visualisation as a performance also shows another advantage of shared feedback. By making the feedback public, it can be used for creative expression. Although one might argue that feedback to augment a performance should be a separate application in itself, the performative nature of this visualisation does have potential learning benefits as well. Many of the participants were motivated to practise because they were going to perform; having public visual feedback during practice may well give some of the reward of performance during in practice making it more motivating. None of the participants showed any concerns about having a shared display where others could see their bowing, in particular Cathy and Eleanor specifically said they did not mind and Andy suggested using it for performances.

There are also challenges to designing shared real-time feedback. Giving feedback about multiple players in a way that participants can make comparisons involves representing more complex information than in the previous study. Given that some people had difficulty using the visual feedback in the previous study it is not surprising that the participants reported that the shared display was very difficult to use while playing from sheet music. Moreover, ensemble practice is more demanding than individual practice making it even more difficult for players to be aware of the feedback system. This is demonstrated by the more extreme cases of participants not being aware of the vibrotactile feedback even though it was the same as that used in the previous study. Therefore, trying to give peripheral shared feedback may not be possible. Instead it may be more appropriate to design the feedback to be used as a central focus. However, this limits its usefulness to situations where ensemble players already know the music well.

Another challenge of shared feedback is where it should be positioned so that all can see it. In this case we had a shared projected display on the wall in front of the players. Again this limits the potential for people to use it with sheet music as it is difficult to switch been the short distance of reading the music and the visualisation which is much further away.

A third challenge with a shared feedback is that all the participants need to have a shared objective when using the feedback. While this may appear to limit the usefulness of shared feedback, it could also be beneficial for participants. If the feedback can be designed in such a way to support analysis and discussion about how the ensemble should be aiming to play, then the requirement of a shared goal may be the catalyst to stimulate that discussion.

6.8.2 Designing Real-time Feedback for Ensemble Playing

There is no single design that will be optimal for every ensemble in every situation. However, the findings from this study do raise some key design decisions that need to be made when designing a shared display for an ensemble and the trade-offs involved in making these decisions. These include the positioning of the display, whether it attracts peripheral versus central focus of attention, how information is displayed, modality and cognitive demands and objective set.

Positioning

The visualisation used in this study was a large projection at the front of the ensemble. This has some advantages in that having a large shared display allows all participants to point to it easily and discuss it. However, it also has disadvantages, it cannot be easily seen in peripheral vision while playing from written music. It also encourages participants to position themselves in a line facing towards the display whereas ensemble players would normally play in a circle or arc so that they can see each other more easily and this change in positioning was observed in this study when the visual feedback was introduced.

There are other design options. One would be to make the display a physical artefact that the ensemble can cluster around rather like a conductor. However, this would not overcome all the difficulties with viewing it in the peripheral vision. Another alternative is that visual feedback could be positioned on the music stand; this would overcome the difficulties with peripheral vision and playing position of the members of the ensemble. However, this raises questions about how the display remains shared. There are two main options for keeping the feedback shared. The first option is to display the whole group's feedback on every stand. The second is to give individual feedback on each music stand, but position it in such a way that the other members of the ensemble can see it. The second option has the advantage that it encourages players to look at one another which is important for ensemble playing and it prioritises personal feedback over group feedback which may make it simpler to read. However, it does separate individuals whereas the original aim of the design in this chapter was to show individual contributions as part of the whole ensemble.

Peripheral versus Central Focus

A key decision that needs to be made is whether the real-time feedback should be designed to be given a peripheral level of attention or should be designed to be the central focus or whether it should be shown to players after they have finished playing or after a practice. Feedback that can be used peripherally has a wider application but it is very hard to design for an ensemble situation. The visualisation in this chapter was not well suited to making participants peripherally aware of their bowing but worked well as a central focus. This is because it showed too high a level of detail and was too unobtrusive. More obtrusive elements could be added to a visualisation design for peripheral awareness, for example flashing lights and other sudden changes to draw the attention at key times. However, care should be taken as findings from the previous chapter show large differences in people's perception of these features. In addition, with a shared display it would also be hard to judge when the appropriate time to draw attention to the display would be as not all members of the ensemble would necessarily need to change their playing. This would be more feasible on displays mounted on the music stand as feedback could be individualised. If the feedback is design for use as the central focus while playing, or is given after players have finished playing, then there is a lot more scope for adding detail and complexity to the visualisation. This additional detail may indeed be necessary because it will need to be more engaging for the player as they will need to feel motivated to do this additional practice either in real-time with well-known pieces or exercises, or to take the time review the feedback after they have finished playing.

Displaying Information

The visual feedback was easily understood by all the participants. Even the more complicated aspects of the visualisation, such as bow speed were commented on by participants, demonstrating that they understood what it was showing. Although the flower shape may initially appear ambiguous in terms of what it is representing, it is still easy to learn how the shape relates to bow use because it changes in real-time in direct response to people's movements. This is an advantage of real-time feedback; there is scope for unusual designs and no explicit labels or explanations are needed in the design so long as there is a direct connection between movement and feedback. On the other hand real-time visual feedback is limited by the constraint that people can only look at one thing at a time and this is why the flower was particularly good because it gave information about length and speed in a single shape.

There is a trade-off to be made between giving detailed feedback which will be applicable to more people, and its ease of use as peripheral feedback. How to balance this trade-off is dependent on the needs of the players and whether the feedback is going to be used peripherally or as the main focus. As already discussed, the visual feedback was too complex to use as peripheral feedback but did work when it was a central focus. Players with a simple learning goal such as long bowing are more likely to find a simple display like the vibrotactile feedback more useful, whereas advanced players who have more complicated goals such as using bowing for expressivity may prefer more complex feedback that enhances expressivity. Publicly visible feedback may be particularly good for these advanced players because it may augment a feeling of performing.

Modality and Cognitive Demands

As discussed in the previous chapter, people vary in their abilities to be aware of and interpret visual and vibrotactile feedback while playing. This study shows that this variation is exacerbated by the demanding nature of ensemble playing. This presents a challenge when designing a shared system to make it accessible to all members of the group. In this study we took a multimodal approach giving feedback in two different ways and this was helpful as some participants used the visual feedback more and some used the vibrotactile feedback more. As well as differences between members of an ensemble using the feedback there will also be changing needs depending on external circumstances. Therefore a flexible multimodal approach may be best, to accommodate varying user needs.

Objectives

Setting an objective for the group which everyone felt was beneficial was key to the successful use of the ShareSense. The design of the feedback needs to support this group decision making. First, for some of the participants, the system did not offer feedback on something that they needed to improve and therefore it was not useful to them. To combat this, the system could either offer more kinds of feedback or it could target the system more carefully at groups of participants who have objectives that match up with those of the feedback. Second, the design of a feedback system needs to help participants to identify problems with their playing and set these as their learning objectives. For example, more detailed feedback could be given when first using the system so that participants can learn about how they are using their body before settling upon a particular objective. This detailed feedback could be given after playing as well as in real-time so that the player is able to reflect on it.

The physical act of calibrating the system could have also been better designed to support the process of setting objectives. The arduous process of the calibration in this study, which involved adjusting and then testing, may have deterred people from changing their objectives. A clearer mapping between the actions involved in the calibration and the goal that is being set would make it easier to understand and might encourage people to tweak and play around with different calibrations. Turning a knob is hard to relate to changing the length of a bowing movement, whereas perhaps recording some example bow movements may be easier for participants to understand.

6.9 Contributions to the Framework of Multimodal Real-time Feedback

In the previous chapter we began to build a framework to understand real-time feedback. The first element of this is *Catching Attention*, and this seems to have been central to how it was used in this study. The study of MuSense showed that there can be big differences between people's ability to notice vibrotactile feedback. This was found to be more so in this study with some participants reporting that they never felt it while playing, whereas others believed they were able to feel it whenever it was triggered. As discussed, ensemble playing is more demanding than individual practice and it appears that these extra demands are the reason for this increased polarisation in the ability to be aware of vibrotactile feedback. This is demonstrated in the way that participants who could not feel the vibrotactile feedback during the ensemble playing could feel it in less demanding tasks like playing scales.

The findings also show that many participants found it very difficult to use the visual feedback while playing in an ensemble. In some ways this is to be expected because the visualisation did not employ flashing or sudden changes in the display which the previous study had shown to be effective for catching attention. Loads on the visual channel come from looking at the music and other players in ensemble playing, or a specific aspect of the design of this visual feedback, such as its position, which makes it impossible to monitor peripherally. Interestingly the one participant (Andy) who confidently reported awareness of the visual feedback found the vibrotactile feedback very difficult to be aware of; whereas the two participants (Cathy and Eleanor) who were confident they always felt the vibrotactile feedback reported serious difficulties noticing the visual feedback. This shows individual differences in sensory perception, which may come from multiple sources. First, there may be inherent differences in the individual's sensory perception such as differences in ability to perceive vibration (Halonen, 1986) or awareness of their peripheral vision (Williams & Andersen, 1997). Second, there may be differences in their approach learning such as having a visual, auditory or kinaesthetic learning style (Fleming, 2006). Third, the difference may come from the way they play music: Andy often improvised, which does not rely on the visual modality but does involve high cognitive load; Eleanor and Cathy both played from sheet music which loads the visual channel as well as being a cognitive load. What is interesting here is that the cognitive load from Andy's playing seems to affect his ability to use the vibrotactile modality but not the visual modality, whereas Eleanor and Cathy's ability to use the vibrotactile modality is not affected by the cognitive load of reading music.

Another finding that is relevant to the framework is that Andy said he found it hard to notice the feedback because he did not have a clear objective. This demonstrates the crossover between the different elements of the framework. More specifically, that not having a clear objective in terms of using the feedback to take *Action* makes the feedback difficult to *Interpret* which results in it not being able to *Catch Attention.* This shows how rather than being four separate stages, the elements of the model are all interdependent. One way to explain this overlap is the Theory of Event Coding (Hommel, Müsseler, Aschersleben, & Prinz, 2001). This asserts that both perceived events and planned actions are coded in the same way, making an overlap between intention and perception. Therefore, it might follow that feedback which is more closely related to the participant's intentions will be easier to perceive. However, this could equally be explained using arguments concerning motivation: if the feedback does not fit their learning objectives then they will not *value* it (Eccles & Wigfield, 2002) which following a *expectancy x value* model of motivation means that they will not be motivated to dedicate effort to trying to be aware of it.

6.10 Conclusions

This study suggests that shared real-time feedback can support ensemble playing and group discussion. However, it is a dependent upon a number of factors that come from the design, the environment and the individuals involved in the ensemble. Ensemble players must have a shared goal that they are all motivated to improve and the feedback must be designed to support this. Ensemble playing is an inherently demanding setting: there are additional cognitive demands in terms of keeping time with the other players and visual demands as players normally monitor their fellow players in their peripheral vision in the absence of a conductor. Therefore feedback should be simply designed and all participants need to be comfortable with the repertoire. Ideally participants should be well matched in terms of level of ability to aid confident discussion about the group playing. Positioning shared visual feedback is also a challenge as players need to be able to switch their gaze between music and feedback; a projection on a wall was not effective, a different approach will be taken in future designs. Calibration also needs to be carefully designed to facilitate setting goals; the design used here did not achieve this as the mapping between calibration and bowing goal was not explicit. A more user-centred and embodied approach is needed for future designs of calibration.

Being a participant observer provided me with many insights into how the ShareSense prototype was being used and how it might be used in future settings. In addition to the interpretive benefits of participation, there are also benefits from a design point of view as it provides a first person experience of using the prototypes that enables empathy with users (Wright & McCarthy, 2008).

Participation has enabled a useful interpretive perspective on the qualitative results from this study and has been informative interpreting findings about the nature of real-time feedback for ensemble playing that are relevant to both theory and design. It has also presented a lot of challenges in the way the study was conducted and has required a great deal of reflexivity about the role of the research in influencing the research and interpreting results. Depending on the nature of the study it may not always be justified or practical to take a participative approach. In our next study which returns to the core user group of violin learners, namely school children, it would not be possible to participate as an equal participant in the research. However, in this study I still actively participate, but in the role of a teacher/helper instead.

The summer school was only a week long, which did not give time for participants to become familiar with the feedback and appropriate it to their own goals. It also prevented prior research with the participants to ensure the design was matched to their learning objectives. In the next study, I investigate real-time feedback in the longer-term with children from a high school orchestra. This study aimed to explore how feedback may be integrated into the process of learning good posture for the violin involving different types of practice such as ensemble practice and individual practice at home. In particular, it aims to discover how children will choose to use real-time feedback when they are given their own systems to practise with.

6.11 A Final Word on Methodology

This study employed a different methodology to the previous study, namely being a participant observer throughout the week. My experiences, the challenges and many self-reflections are reported in depth in an article that was published in CHI (Johnson, Rogers, van der Linden, & Bianchi-Berthouze, 2012). Here, I highlight some of the most salient using extracts from my diary. Time and tiredness were foremost in my mind:

"Full days meant that as the week continues, I and others around me, began to become tired. I also began to suffer from pain and numbness in my shoulders (due to my own bad technique). As I started to take pain-killers myself I began to notice the odd packet of pills on other people's tables at lunch. Then, one participant asked to miss a session due to back pain. No-one wanted to give in to these physical demands though – a sign of how much people love playing at the school."

Time is not of equal value in every context; time at the school was much more valuable than it would be in many other places. Taking part in the study asked participants to give up time which could be spent on many other activities. I came to understand this through participating in the summer school. It meant that I interpreted missed sessions differently because I understood the contextual reasons for them. It also meant I valued the compliment when one participant chose to attend every session and saying how useful she found it.

Another tension I was aware of throughout the study was between being too prescriptive and giving enough information to make participants feel confident using ShareSense. It is difficult to explain how to use something without strongly encouraging its use and specifying how it should be used. At the same time it was important not to enforce a particular learning objective on the participants because the aim of the study was to investigate how ShareSense could facilitate players in achieving their own learning objectives. Often the researcher comes across as authoritative, even omniscient within the domain of the study. However, I felt quite different:

"I did not wish to take on the role of authoritative researcher as I was interested in what participants chose to do with the feedback themselves rather than instructing them or enforcing any particular task upon them. Nor was I in any position to take on the role of a teacher, most participants were considerably better players than myself. In some ways this felt uncomfortable as I do not think this fitted with their expectations, however, I felt that by giving up the authoritative position I became a more equal participant."

In doing so I was able to encourage the participants to question aspects of the study and take the lead in the practice more.

Confidence in myself and the prototypes also played an important role in how I interacted with participants:

"My rapport with Cathy was different from the others because I believed that I and the technology could help her improve her bowing and because she believed it too. We also had a lot of previous contact via email and she played a style of music I was more familiar with. All this meant I felt I could give some advice and point out areas where she could improve or should try and listen to the prototype. This sometimes reversed when she pointed out that I was not making as large shapes on the visual display as she was and this motivated me to try harder to use longer bows in my own playing."

"I felt less confident doing this with the other participants because I was too conscious of my own inadequacies. Instead I tried to be a good example of self-analysis by saying aloud whether I had managed to trigger the feedback and what I thought that meant."

My reflections show how confidence and a shared belief in the potential success of the study can create a positive rapport between the researcher and participant. This in turn caused me to have the confidence to contribute in a more hands-on way.

Chapter 7 – Buzzy Jacket and Twinkly Lights: Studying Real-time Visual and Vibrotactile Feedback in Different Practice Settings

7.1 Introduction

The research presented so far has focused on investigating factors affecting the use of feedback for practicing playing the violin in the short-term. Here, we consider whether there may be other factors which only become apparent or important when practicing in the longer term. The study presented in this chapter considers how real-time feedback is used as part of the long-term process of learning to play the violin through individual practice at home and ensemble practice. The aim is to study how the two types of real-time feedback already investigated in this thesis - visual and vibrotactile - can be incorporated into a longer term programme of learning. It addresses the third research question from section 1.3:

How can real-time feedback systems fit into the extended process of learning to play the violin through various forms of practice?

The research reported in the previous two studies found that real-time vibrotactile and visual feedback can be effective for teaching the technique of long bowing – for both individual and group practice - provided that there are not excessive demands on the player's attention. An overarching finding was that perception of real-time feedback is reduced under high cognitive load. Given this limiting factor, the focus in the final study was to explore how and when visual and vibrotactile feedback can help learners when practising in different contexts. Its role in this context is seen primarily as providing increased awareness of how they are using their body, and secondly, its motivating capacity.

This study turns attention to the largest potential user group for real-time feedback, which is children learning the violin. The previous two studies with adults have been useful because adults may be more easily able to express what it is like to use the feedback. However, it is important to also test the feedback with child learners because most people learning the violin are children, in particular many start learning at school.

A study was designed to look at how children learn in different settings over a long period of time. Specifically, it explores how five players aged 11-13 years, from a high school orchestra practice at home and at school and how the introduction of feedback technology impacts upon their routines. Children of this age are likely to be distracted by other activities and may give up playing as a consequence. A goal was to see if visual and vibrotactile real-time feedback could be designed to make practicing more enjoyable while also providing cues for where they could improve their playing. Hence, could technology interventions make learning to play more enjoyable? Most musicians start learning as children and positive childhood musical experiences have been shown to play a key role in shaping successful musicians (Sloboda, 2004), (Lehmann, Sloboda, & Woody, 2007).

Based on the results arising from the research so far, a framework of how people can use different types of real-time feedback to improve their playing has been developed. It considers how best to design technology interventions that can be used to help learners with the processes of catching attention, taking in information, interpreting and taking action. The research findings also suggest that each individual interacts with the world slightly differently, for example, one person's attention may be more easily diverted than another. How feedback is experienced and utilised is controlled by each person's unique set of characteristics, for example, someone whose attention is easily diverted may dislike feedback which is too salient because it dominates their attention. Individual differences were also identified in people's responses to different modalities of feedback. Having observed these individual differences in the previous studies, this study also aims to explore whether preferences for particular modalities and styles of feedback are related to differences in approach to music learning.

The methodological approach adopted was again an in-the-wild study, where data was collected from the participant's home practices as well as nine ensemble practices during school time run over an eleven week period. The design and implementation of real-time feedback and how it was presented to the children in the form of two systems - a *buzzy jacket* and *twinkly lights* - is described in the next section. An analysis of how the two systems were used by children practicing at home and at school is then presented, followed by results from interviews conducted at the end of the study.

7.2 Design Rationale

Buzzy Jacket was designed as a hoody type of jacket to be put on and worn whenever the child was practicing playing the violin. Hence it needed to be an item of clothing they thought fashionable, was comfortable to wear and which they owned. Much thought went into selecting the kind of clothing that fitted these criteria. Embedded in the jacket were a set of sensors and actuators intended to give feedback about the posture of the upper body. The reason for focusing on posture was because this was a particularly relevant goal for the participants in this study. Initial observations of the participants playing in the school orchestra before the study began showed that they had a tendency to slip into bad posture such as hunching over and letting their violin slip down which had a detrimental effect on their playing, and this was an issue that the conductor of the school orchestra wanted to address. Hence, the focus of this study was on how the feedback could increase their awareness of their posture, and in turn improve their playing – rather than focusing on improving a specific technique, such as long bowing – which was the focus of the previous studies. The reason for this shift of focus is that it was considered to have more general utility for helping and motivating children to play, as good posture is an important goal for all players irrespective of ability or the particular repertoire they are playing.

Twinkly Lights was designed as the visual feedback component intended to give real-time feedback about each child's foot position. Again, this shift of focus to using visual feedback to provide immediate feedback about an aspect of an individual's playing in a group rather than a group's collaborative performance was intended to provide information that could be attended to peripherally and was not too much of a cognitive overload. To this end, a footboard attached to twinkly lights was developed that switched on when the feet were not correctly positioned. This provided a contrasting modality of feedback to the vibrotactile jacket which allowed for different modalities to be discussed in the interviews. Both prototypes could be used together as well as separately to give a multimodal experience. Each of the prototypes is described in more detail.

7.2.1 Buzzy Jacket

The Buzzy Jacket was designed with vibrators placed on the ribs and elbow which would vibrate if the participant was slouching or holding the violin incorrectly (respectively).

Sensing

The vibrotactile jacket used two accelerometers connected to an Arduino pro mini. One accelerometer was positioned on the inside of the forearm and the other on the left shoulder blade as shown in the diagram (Figure 7-1). In this design the accelerometers are being used as inclinometers (measuring angle of tilt) to measure details of the player's posture. This is only possible because the parts of the body the sensors are placed on will not be accelerating so it can be assumed that the accelerometer is mainly measuring gravitational force rather than the movement. The accelerometer on the forearm was used to measure the angle and orientation of the arm and from this it could be inferred whether the player was holding their violin correctly (Figure 7-2 shows how this can be done). The accelerometer on the shoulder blade was designed to pick up on whether the shoulders were straight or hunched. Interpreting the measurements from the accelerometers relies on having recorded an ideal position to compare the readings to. This was done through a calibration process.

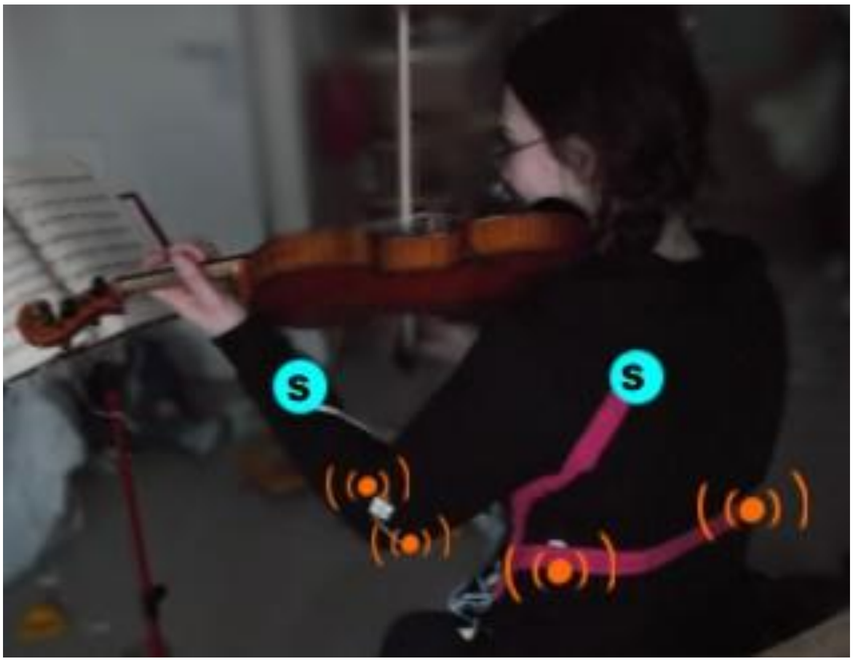

Figure 7-1: Photograph of the Buzzy Jacket. (S) Shows position of accelerometers. ((o)) Shows the position of the vibrators.

Vibrator Position	**Meaning**	**Priority**
Left and right Ribs: activated together	Straighten up back: triggered by slouching	Top of hierarchy
Elbow, near funny bone	Lift up violin: triggered by drooping violin	Bottom of hierarchy
Left side of Elbow	Move elbow under violin: triggered by elbow moving to side	Middle of hierarchy

Table 7-1: Vibrators positions and corresponding instructions

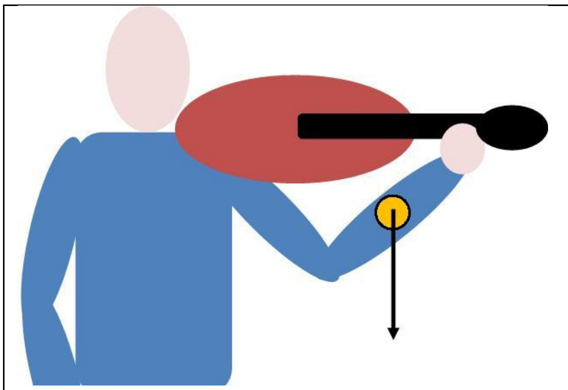	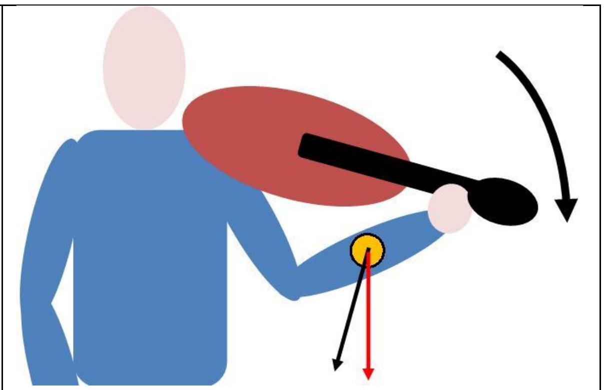
Before using the feedback the player sets their ideal position and the reading from the accelerometer is saved as a reference (black arrow).	Once the player starts playing the current (red arrow) measurement from the accelerometer is compared to the reference measurement (black arrow) to infer if the player has let their violin drop.

Figure 7-2: Diagram of how angle of arm can be measured using an accelerometer

Vibrotactile Feedback

Four 10mm coin vibrators (Precision Microdrives, 2014) were used to give feedback about posture. These were positioned as shown in Figure 7-1 and were used to indicate to the child that they needed to improve their posture using the corrective action summarised in Table 7-1. Four vibrators were used to give feedback about posture. The two positioned on the ribs were set to come on when the accelerometer on the shoulder detected a hunched posture. The vibrator underneath the elbow came on when the accelerometer on the arm detected that the violin arm was drooping. The vibrator on the side of the elbow came on when the same sensor detected that the elbow was twisted out to the side. Three metal popper fasteners were used as switches so that participants could control which vibrators were active. If they chose to have several vibrators switched on at once, then there was a hierarchy about which one would be activated, so that only one set of vibrators could be on at a time – this hierarchy is shown in Table 7-1. This was chosen to avoid the participants feeling overloaded by multiple vibrators activating at the same time.

Slouching was sensed by detecting movement on the ribs and elbow: this would result in the rib or elbow vibrating separately. Previous work with the MusicJacket (van der Linden, Johnson, Bird, Rogers, & Schoonderwaldt, 2011) had shown some participants had needed to bring their elbow underneath the violin before pushing the violin up, but the feedback on the MusicJacket had not been able to tell them this. Hence, it was decided here to add another motor to indicate to the player if their elbow came out too far to the side.

The Buzzy Jacket differed from the MusicJacket in another key way: the minimal sensor set-up meant it was much more affordable. This enabled five jackets to be built at a cost of approximately £70 each (whereas the MusicJacket cost £10,000 to build), one for each participant to take home. The jackets were specially made for each individual, using jackets in their size and favourite colours.

Each Buzzy jacket was powered by three AAA batteries protected by a 0.5A fuse. Initially a small lithium polymer battery was used because they are lighter (for the same capacity). However, these had two drawbacks: first, the children could not replace or recharge them themselves; second, lithium polymer batteries can overheat explosively when short circuited and given the homemade nature of the Buzzy Jacket and the fact the batteries would be worn close to the body it was considered safest to choose the safer AAA alternative.

Calibration

An initial calibration of the ideal playing position was required in order for the sensors to be able to judge whether the future playing positions were correct or not. It took a second design iteration to design a way for the children to be able to calibrate the jacket easily by themselves.

To calibrate the final version of the jacket, the wearer had to press a button in the middle of the jacket near the zip. The jacket would then buzz five times to give them time to get into the ideal position, after the fifth buzz their current arm and back angles were recorded as the calibration. This worked well for gaining a reference to compare their future arm and back angles to.

7.2.2 Footboard and Lights

At the start of the user study, baseline data was collected about how participants normally practice. At this point it was noticed that participants had difficulty remembering to keep their feet firmly on the floor when playing. This issue needed to be addressed before other postural improvements could be made. The Footboard and Twinkly Lights (Figure 7-3) described in this section were built to nudge the participant to put their feet back on the floor if they lifted them while playing.

To save on cost, the design combined two off-the-shelf items. On the sensor side it used the contents of an electronics project kit for school children to make a light sensor circuit. This contained the necessary PCB (Printed Circuit Board) ready-made and was adapted by adding further LDRs (Light Dependent Resistors) in parallel and adjusting the value of the resistors in the circuit. In this circuit the LDRs are being used as light sensors. This was then connected up to a set of battery powered LED fairy lights. These had a built in digital effects setting which made the lights flash in different combinations. The circuitry in the fairy lights was adapted and combined with the LDR circuit so that the lights would come on and create a colourful display when the LDRs were in the presence of light.

The LDRs were set into an A3 size piece of card onto which the children's footprints had been carefully drawn. There were two LDRs for each footprint, one under the ball of the foot and one under the heel. When the children's feet were in the correct position the LDRs would all be in darkness so the fairy lights would stay switched off (Figure 7-3 left). If the children moved their feet off the footprints the LDRs would be in the light causing the fairy lights to be switched on (Figure 7-3 right). Each footboard had the child's unique footprints drawn onto a white background. The children were encouraged to decorate their footboard at home - one child did this.

Real-time ambient visual feedback was chosen for the Footboard and Twinkly Lights because findings from the design of the MusicJacket (van der Linden, Johnson, Bird, Rogers, & Schoonderwaldt, 2011) suggest additional vibrotactile feedback would be overwhelming. A multimodal approach on the other hand offered the opportunity to give feedback about the foot position and upper body posture at the same time, with the contrast in modality allowing the user to be able to make a distinction between the two. The studies in the previous two chapters show that visual feedback could be used even while reading music and was best positioned on the music stand in order to be visible in the player's peripheral vision. Therefore, the Twinkly Lights were wound around the music on the stand. When the feet were positioned correctly the lights would be switched off to prevent them distracting the students from their playing. When the students made a mistake (i.e. moved their feet) the feedback was activated in order to catch their attention and allow them to correct their foot position. The two previous studies with visual feedback had shown that lights coming on was a more salient signal than moving or changing colour. The digital effects that were built into the lights meant they made a colourful moving display while switched on in order to be particularly attention grabbing. However, this raised the question of whether the children would find it counter-intuitive to have an attractive display used as negative feedback.

Figure 7-3: Footboard and Twinkly Lights – the lights are activated when the child moves her feet off the footprints

7.3 Methodology

An in-the-wild study was conducted with students from a state school orchestra. Regular meetings were held with them for group practices in a school room which were videoed. Additionally participants were also given a video camera to collect data about their home practice. To begin with, the participants were given feedback about their posture while playing, initially verbally and later using the Buzzy Jacket and Footboard and Twinkly Lights. At the end of the 14 weeks study, the participants were interviewed about their experience using this multimodal programme of feedback. The multiple sets of data (from home practice, group practice and interviews) were intended to complement each other giving both the perspective of an observer (video and researcher observations) and the participant (interviews). They were also used to triangulate the findings, using video evidence to confirm ideas put forward by participants in the interviews.

The researcher conducting this study underwent an enhanced disclosure CRB check carried out by UCL because she was working with children. Participants were recruited by the music teachers at the school. Before the study commenced potential participants had the project explained to them in the presence of their teachers and an early prototype was demonstrated to them by the researcher. Consent forms and information sheets were then sent to the parents of those who expressed an interest in taking part. This was organised by the school and an information sheet about the study was provided by the researcher (information sheets can be found in appendix F.7). Participants only took part once written informed consent had been received from their parents. During the first session the study was explained to the participants again and they had the opportunity to ask questions, the participants then signed a consent form which can be found in appendix F.4. During the first session the orchestra conductor was also present as a representative from the school to oversee how the study sessions were run.

7.3.1 Participants

Five high school pupils (Alice, Bea, Carl, Daphne and Esme) participated in the study, three played in the third violin section of the school orchestra and two in the viola section. Pseudonyms have been used to protect anonymity. The way the parts were arranged in the orchestra meant that the viola section and third violin section both played the same part although sometimes an octave apart. The participants were aged between 11 and 13 years and comprised four girls and one boy. There was a variety of abilities spread between grade 1 and grade 5. They were friendly with one another and much of the session was characterised by chatter and banter. There was a competitive streak between them as well as a desire to help one another; this is described in more detail in the results section as it affected how the feedback was received in the group. Not every child was able to attend every session as demands of other school activities such as choir practices and the French exchange drew them away. The average attendance was between three and four participants at each session. Similarly, once the feedback equipment had been handed out they did not always manage to remember to bring it to every session. The exigencies of everyday life, like forgetting to take things to school are an unavoidable feature of an in-the-wild study.

7.3.2 Ensemble Practice Sessions

There were nine ensemble practice sessions. The sessions were weekly with breaks for half term and Easter holidays. Each session lasted an hour and fitted between the end of school and the start of orchestra practice at five o'clock. During the first six sessions the participants played together practicing the pieces for an upcoming inter-schools concert. The concert was held after the sixth session and in the remaining three sessions participants played other repertoire, such as rounds and duets.

I led the group by playing violin, conducting or singing the cues from other instruments so that participants could practice their entries. Although posture was the focus for the technology intervention, much time in these sessions was given over to learning the notes and rhythm as well. The motivation for the sessions from the students' perspective was to support them in learning the pieces for the concert while helping them to play with good posture. This context was chosen because it made the sessions more directly useful to the children and was a more naturalistic representation of music learning than only focusing on posture.

The reason for these sessions from a methodological perspective was six-fold. First, working with the students in these sessions contributed to a better understanding of their needs which informed the final design for the prototypes. Second, I could lay the foundations for using the prototypes by teaching them about how to sit and stand so that they were aware of the importance of posture and what to aim for. Third, once the Footboard and Twinkly Lights and Buzzy Jacket were given out, the weekly meeting allowed me to monitor how they were working and give out new batteries where necessary. Fourth, the sessions gave an opportunity to observe and video record the prototypes being used, whereas the recording of home practice were made by the students themselves and would be dependent upon whether they had time and were able to do this. Fifth, meeting the students every week in a sociable practice setting meant that I could get to know more about them and build up a relationship with them. Lastly, ensemble practice is a unique type of practice, very different from individual practice, and it would not be possible to infer from a study of home practice how prototypes would be used in an ensemble. It is, therefore, important to study the role of real-time feedback in both ensemble and individual practice settings.

Session content

Most sessions began with a warm-up and discussion of good posture. Participants were asked to shake out the tension in their body and then stand up straight by imagining their "head was a ball balancing on a fountain". After this they would practise sitting correctly on the front of the seat with their weight slightly forward, feet firmly on the ground, back straight and neck in line. They would then practise by playing a game where they had to stand up as quickly as possible when I clapped my hands and if their feet were well positioned and their weight correctly distributed they would be able to stand without losing balance.

Next, time was given over to practicing pieces. Difficult parts were played slowly and then sped up. Difficult rhythms were clapped or played on a single note before playing it at the correct pitch. Sometimes I used recordings of a piece for participants to play in time with. Over the sessions participants became more familiar with the orchestral pieces and were able to play them faster and more proficiently.

Figure 7-4 gives an overview of what happened in each session. Not all the children were present at each group practice and in some practices there were also adults from the orchestra present. In the first session the conductor sat with us throughout. In sessions 2, 4 and 6 a helper from the orchestra attended for the last 15 minutes of the session. When the orchestra helper was present the focus was more strongly about learning the orchestra pieces and less about posture or the feedback. Furthermore, the precise content of the session changed over time, depending on how close the concert was, which children were present and what feedback was being used. Before the concert only orchestra pieces were practised, while after the concert other pieces were played to offer variety and to keep the children interested. On the week when only one child came we worked specifically on his own pieces. On the weeks where the prototypes were given out for the first time, considerable time was given to explaining how to use them. In particular the calibration process for the Buzzy Jacket was explained and practised in detail in week 9 to ensure that the participants understood how to use the jacket and that feedback was coming on in the correct place.

Verbal feedback

During the practice, I gave verbal feedback about their playing, as a teacher might; commenting on how well they were holding their instrument up and whether they had moved their feet while playing. Generally, the verbal feedback was directed to the group as a whole rather than singling out individuals. However, on some occasions I did direct comments to specific individuals, often in response to questioning from students. I tried to balance praise and criticism: praise to keep the students feeling positive; criticism to make them aware of what they needed to do to improve their posture and to prevent them from being surprised when the technology pointed out their mistakes. After the Buzzy Jacket and Twinkly Lights had been introduced I gave less direct feedback about posture, instead asking whether the participants had felt/seen any of the real-time feedback. As well as verbal feedback about posture I gave other feedback about playing, for example, pointing out that there were E flats in the key signature or that different participants were playing at different speeds.

Data collection

The sessions were videoed from two different angles. Having two views was important because often the view would become obstructed during a session by people standing in front of the camera. Videos of the sessions, plus my notes written before and after the sessions are the basis for descriptions given later in this chapter about how each prototype was received and used by the students during the group practices. Videos were chosen not only as a record, but to allow me to go back and view the sessions from a different perspective. While I was running the sessions I was thinking about teaching. When watching the videos later I could think more as a distanced observer. The written notes and my memories of the sessions are a contrast to this; they record my immediate views while looking at events from a teaching perspective.

7.3.3 Home Practice and data collection

One of the aims of this study was to investigate how real-time feedback could be used by children in self-directed practice at home - whether they would choose to use it at home and what for. Collecting data about home practice is difficult. A researcher cannot fix a time with a participant and come and observe or film it naturally occurring because to do so would be to force practice to occur at that time. Practice diaries could be used, but these place a burden upon the participants, rather like homework, to write down what they did. For these reasons I asked the participants to film themselves. Initially, it was thought participants may be able to do this using camera-phones. However, most the children did not have mobile phones, and those that did had ones designed only for making phone calls. Instead each participant was given a low cost video camera. Most participants did try to video themselves practicing. However, the resulting films were not the fly-on-the-wall views into home practice as originally envisioned. Instead they took on a more performative character. Participants were selective about what they recorded. Participants only filmed themselves playing their pieces, not scales or exercises, and they nearly always played them all the way through without making mistakes. When mistakes were made the video would be stopped and the next video would show them playing the piece again from the start. The film numbering on the camera revealed that some videos had been deleted by participants. It was clear that participants saw playing for the camera more as a performance than a practice and so were censoring themselves. The performance aspect could also be seen in the way one participant was very proud to show the others the videos of him playing difficult pieces. Although these videos do not represent everyday practice, they do give interesting insights into what the children were willing to share with us.

There were also two technical issues which caused difficulties with home videoing. Firstly, the low battery life in these cameras meant that the films did not always record successfully. New batteries were handed out each week to help combat this. Secondly, the participants generally practised in small rooms meaning the cameras could not be positioned far from the participant so the participant's body was not fully in shot; often only their arms and torso were in view. Participants' parents and sibling often got involved with making these films by holding camera while they played. One participant also spontaneously chose to write about her practice with the footboard and later with the jacket. The contents of these written reports seemed to give a different perspective on use to the performances for the camera.

Study Timeline

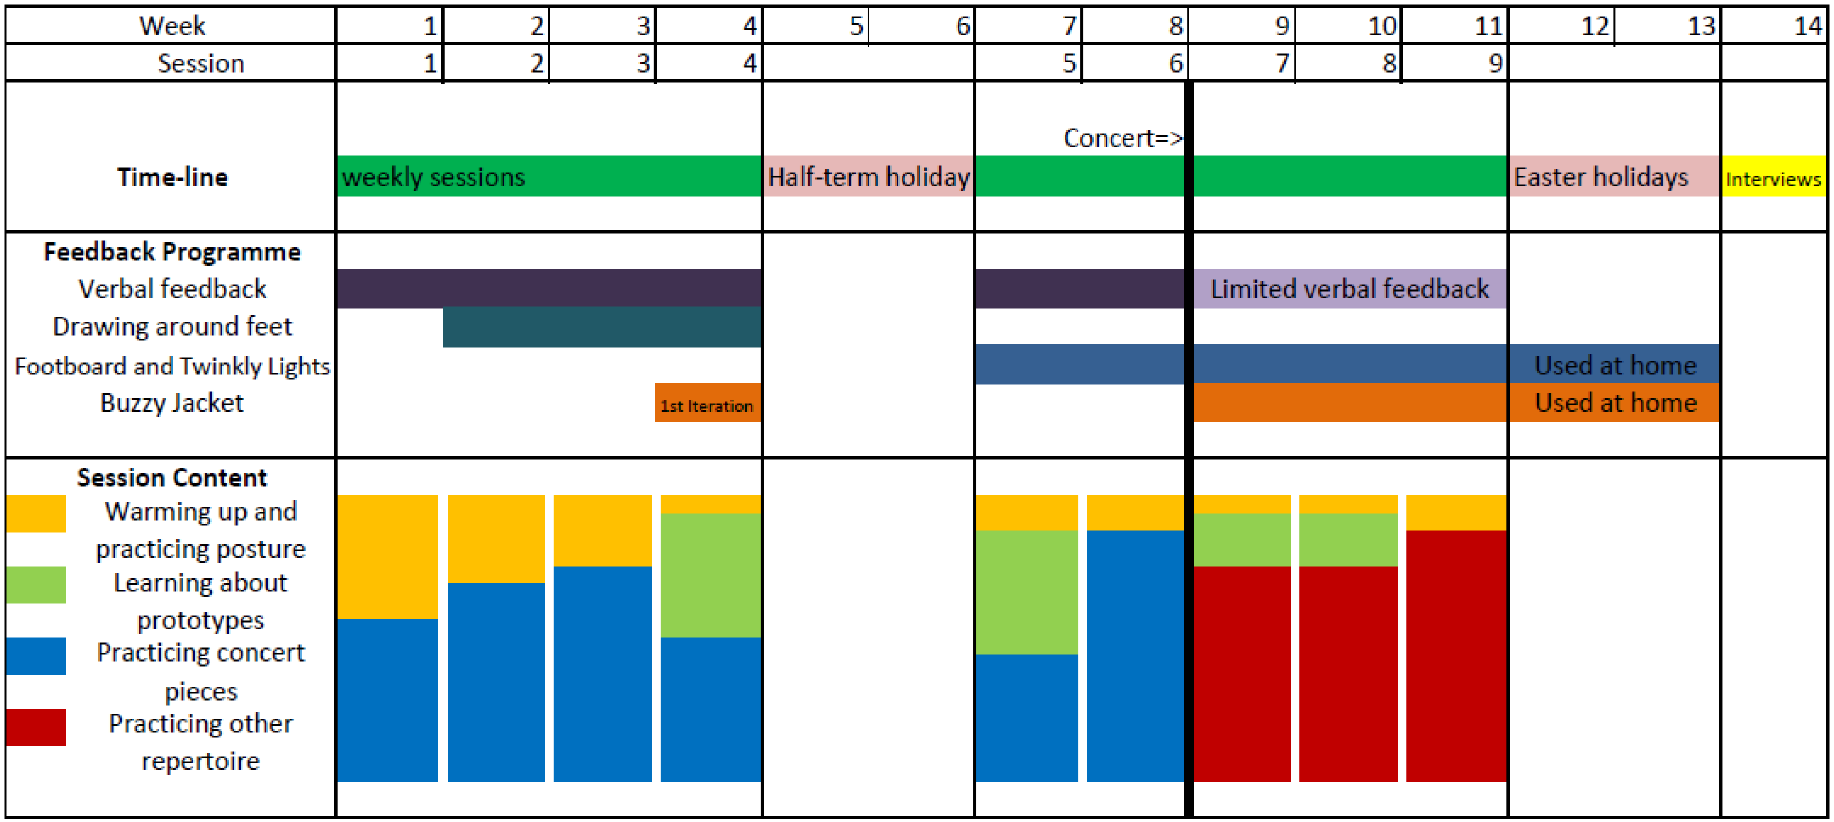

Figure 7-4: Overview of what happened in each ensemble practice session. The Time-line shows which weeks we met for group practices in relation to the school holidays and the inter-schools concert. The Feedback Programme shows the weeks when participants had access to different types of feedback. The Session Content section shows what the proportion of each group session what was given over to particular activities. For example in the first session we had a long warm-up and discussion of posture, but by the 7th session only a small proportion of time was given to warming-up.

7.3.4 Interviews

At the end of the study participants were interviewed. Three of the interviews were with individual participants and one was conducted as a pair with Daphne and Esme who were close friends. The choice to interview the friends together was an opportunistic decision. The interviews were semi-structured, topics for questioning were pre-planned. The interview plan can be found in Appendix D. It covers four topics: practice and how using the feedback impacted on this; comparing the feedback modalities (this section uses a smileyometer see below); general approach to music playing and designing a new feedback system. These lines of questioning were chosen for specific reasons. The section on practice was intended to understand from the participant's point of view whether and in what ways the prototypes were helpful in practice and why. The second two topics, comparing modalities and general approach to music playing, aimed to follow up questions raised in previous studies which showed strong individual differences in their opinions about different modalities of feedback. The questions in these two topics explored whether preferences for a particular modality were also reflected in other elements of their approach to music, for example, whether they preferred playing from memory or reading music or how easy they found the visualisation exercises. The final topic, designing a new feedback system, was chosen to remove the constraints of talking purely about the prototypes used in the study and give the students the opportunity to contribute their own ideas.

In some of the interview questions participants were asked to rate aspects of the two systems on a smileyometer (Read & MacFarlane, 2006) shown below (Figure 7-5). Participants showed their opinion by decorating the faces or adding a nose.

Very difficult Very easy

Figure 7-5: Smileyometer

Researcher Role

The previous studies have demonstrated the importance of the researcher's role when conducting in-the-wild studies. As the researcher in this study, I organised and led all the ensemble practices with the five participants. I also attended the weekly school orchestra rehearsals as helper playing with the third violin section. My role was dictated strongly by the study setting and the participants. As an adult working with children in a school it was natural that I should act like a teacher during the ensemble practices (while making it clear that I was not a qualified teacher). The students automatically treated me in this way and the conductor and music teachers treated me as a colleague around the students. The role of teacher involved giving feedback and comments about playing, answering student queries about music, leading the ensemble either by playing or conducting, managing behaviour and keeping them focused.

In this role, I was able to set-up the foundations for the study by bringing to the students' attention the importance of posture and teaching them what they should be aiming for. When giving them feedback I was able to highlight areas where they could improve their posture so that later when they were given the Buzzy Jacket and Twinkly Lights they could see how they could use them. This meant that student engagement with the prototypes was influenced by how I presented the issue of posture to them and their respect for my expertise and authority.

Being actively involved in the students' learning also allowed me to recognise their needs as they arose and respond to these by dynamically adjusting the plan for the study and the design of prototypes. The biggest change to the study was the addition of a new prototype for foot positioning. On a minor level, the content of the group sessions was adjusted every week to address particular problems the students brought with them.

The role of teacher also gives a unique perspective for interpreting the data. It made me consider the difficulties of explaining physical movement and discover where verbal feedback works and where it does not. When one approach did not work I was forced to think of other approaches, for example trying different mental visualisations and physical props. I became more conscious of the costs of different approaches in terms of how much they distract the group or use up time. I also learnt about the different strengths and weaknesses of the children and their likes and dislikes which gave a useful perspective on how they used the technology and their interview responses.

To summarise, this chapter focuses on how a programme of feedback given over many weeks can help children practice. It is important to understand that in this study the programme of learning was not only made up of real-time feedback technology but also my presence as researcher in the field. I gave all the verbal feedback and taught them how they should think about posture as they used the prototypes. My relationship with the students and opinions about what they needed to work on shaped the design of the prototypes and what we did with them in the ensemble practices. My experience working with them has then gone on to influence how I interpret the data we collected together. The next section describes the findings from this study.

7.4 Findings

The findings from this study draw on data collected during the group sessions (my observations as teacher along with later analysis of the videos), the videos the students took of their home practice and their responses in the interviews. The findings are split into two sections. The first describes what happened in the group sessions as the two systems were introduced. The second section mainly discusses the interview responses, also drawing on the students' behaviour in the group sessions and their home practice videos.

7.4.1 Before Feedback

The learning programme started in the first session with a discussion of what is good posture and why it was important. The students responded to this by asking questions about why I had chosen to work on posture. I gave them reasons involving health, improving sound and looking professional. The orchestra conductor also came to this first group session and supported my arguments giving them extra credibility. By the end the students had accepted that it was worth trying to improve their posture.

The students were quick to learn what posture they should be aiming for. However, once they started to practise the concert pieces their feet would move out of place and their violin would gradually droop downwards. I gave verbal feedback after each piece to point out that their feet had moved or their violin was not held up. This was ineffective because while they were able to correct it between pieces but they would quickly lapse once playing commenced. Foot position was difficult in particular because they were unaware their feet had moved.

Drawing around their feet showed the students how much they were moving their feet over the course of a piece. The effect was that the participants moved their feet much less than previously. Whereas in the first session I had to remind them about their foot position many times, in the second session I did not need to do this. This was a surprising result as they had been unable to do this before and they were not able to look at their feet and see their drawings whilst playing. It seems that by making this external representation of their feet they were able to be more aware of how they were positioned.

7.4.2 Footboard and Twinkly Lights

Although simply drawing around the feet had proved effective for encouraging good foot position, we were interested to find out whether the more interactive approach might enhance and sustain the effect. In session five the Footboard and Twinkly Lights were introduced to the students. The video data showed that the footboard and lights revitalised the children's interest in improving their foot position. It did this by making foot position a more public matter. In particular, it became a focus for the children's competitive interaction. If one child played a rhythm slightly wrong or played out of time the other children pointed it out – in part to be helpful and in part to demonstrate that they knew better how it should be played. This often led to a certain amount of joshing and one-up-manship during sessions. The Twinkly Lights encouraged this behaviour. For example, the children spontaneously started to suggest ways that the Twinkly Lights could be used competitively, as the following conversation illustrates:

Alice: *Can we have a competition and if your lights go on then you're out and then if your lights don't go on then that means you win.*

Carl: *No no. We like have a maximum of ten points and if we get through one whole piece without our lights going on then we [inaudible].*

Bea: *You have ten –*

Alice: *And then every time your lights go on you minus one.*

Drawing around the feet had not stimulated the same type of comments. Comments associated with the drawing were about foot size and whether they should wear shoes. It seems that monitoring feet using the drawings was more of a personal matter, but with the lights it became a shared focus. This meant that even though the lights were not always effective at catching participant's attention (for example Alice played all the way through a piece without noticing them) other participants would now point out someone's lights came on, for example Daphne exclaiming *"Carl, feet! Alice, feet!".* This shows how a publically visible feedback in a group setting can have the knock on effect of encouraging verbal feedback from peers. These kind of comments appeared in sessions five and six (the first two sessions with the lights) but did not persist into the later sessions. This may in part be due to the novelty having worn off and the stronger emphasis which was given to using the Buzzy Jackets in the later sessions.

Nevertheless, even in later sessions, the Footboard and Twinkly Lights was effective at encouraging good foot position. This can be seen by comparing participants playing using the footboard with the sessions where they forgot to bring it. For example, in session eight Bea forgot to bring her footboard and lights and as the left picture in Figure 7-6 shows her feet are crossed and remained like this for the majority of the session. The following session Bea remembered the Footboard and Twinkly Lights and her foot position was much improved (Figure 7-6 left).

Discussion

The findings suggest that both engaging the children in drawing their feet on the footboard and providing real-time visual feedback are valuable and complementary methods in helping students to improve their foot position. Introducing one after the other was able to sustain interest in what otherwise would be quite an unexciting element of technique. In particular, it demonstrates the value of interactive visible feedback as a way of making something like foot position a more public issue which is potentially advantageous in a group like this which is competitive in a friendly manner.

Figure 7-6: Comparing foot position between session 8 when Bea forgot her Twinkly Lights (left) and session 9 when she remembered them (right).

7.4.3 Buzzy Jacket

In contrast to the footboard and lights, the vibrotactile feedback in Buzzy Jacket did not stimulate many peer to peer comments during the ensemble practice. Instead conversations about the jackets were mainly between the students and myself, for example I would ask them if they had felt it in the last piece they played and they would reply when and how many times. The vibrotactile feedback could only be felt by the individual student and so did not make sitting up straight and holding the violin into a public matter as the lights had done for foot position. The students did not monitor each other's posture as they had done with their feet after the lights had been introduced – posture remained a personal goal which students worked towards individually with occasional advice from me in the role of teacher.

One aspect of the Buzzy Jacket which turned out to be important was the way the students had to calibrate it themselves. This meant that the students had to choose what their correct posture was. From the very start of the study we had worked towards them being able to do this through our discussions of what good posture meant, and through exercises and visualisations. By the time the jacket was introduced they could all get into a good seated posture and hold their instruments up well. The Buzzy Jacket was their next step in their programme of learning; it was going to help them maintain this posture while playing. To calibrate the jacket they pressed a button and then got into their good posture position and this was the ‘ideal’ which the jacket would measure them against. What was interesting about the calibration process was that the students could use it to take control of their learning. For example, during the first session with the Buzzy Jacket, Carl recalibrated three times during the practice in order to get the level of strictness he was wanted. He felt the initial calibration was too lax because it did not vibrate at all while he played his piece so he recalibrated higher. At this point he realized he hadn’t turned the feedback on and that it had not been a calibration issue after all. After playing with the very high calibration he felt that the feedback was pushing him uncomfortably high so he recalibrated a little lower. After playing two pieces he chose to recalibrate lower again to give himself the flexibility to sway while playing.

Discussion

Making choices about calibration are important because one potential drawback of giving feedback on posture is that it might make a player keep their body rigid, so recalibrating to accommodate movement while playing is healthy. Moreover, by choosing when he wanted strict feedback and when he did not Carl was taking more responsibility for his learning than simply doing things because he was told to.

In ensemble practices it was harder to ascertain how children were using calibration because it was not discussed directly. They all calibrated the Buzzy Jacket independently during the ensemble sessions and reported feeling comfortable using the feedback and that they were able to respond to it showing that they were able to calibrate it successfully in a position that was challenging but achievable.

After the Buzzy Jacket was introduced I also noticed a subtle change in my role as teacher. Previously I had been giving feedback about their posture and violin hold to make them more aware and inform them how to improve. Now that they had the vibrotactile feedback I thought that they must already be aware that they were not holding the violin up and felt less comfortable giving feedback in the way I had done before because I did not want to nag. Instead I would check that they had felt the feedback and that they knew what to do respond to it. If it was a reoccurring problem I might check how it was calibrated as well. These lines of questioning do encourage the students to improve their posture but they are different to simply telling them to sit up straight or hold up their violin. This questioning makes sure that the student is aware of how they can improve but places the choice to improve with the student. Other teachers with different teaching styles may not respond to using vibrotactile feedback in the same way.

7.4.4 Interviews

The participants had many things in common with each other in their responses to the interviews but also individual stories when it came to their opinions about the feedback. The participant's responses to the interview questions are summarised in Table 7-2 and Table 7-3. Each column shows the responses of an individual participant, each row signifies a question which is described in the left hand column. In Table 7-2 the results are given in two columns for each participant; the left column gives their ratings for using the lights as feedback and the right gives the ratings of using the vibrotactile feedback so that comparisons can be drawn between the feedback types. The ratings were given by participants on a smileyometer but have been abbreviated to a number with 5 being the happiest face and 1 being the saddest face.

Participant:	**Alice**		**Bea**		**Carl**		**Daphne**		**Esme**	
Type of Feedback:	**Light**	**Vibro**	**Light**	**Vibro**	**Light**	**Vibro**	**Light**	**Vibro**	**Light**	**Vibro**
Noticing	3	5	4	4	4	3	5	4	5	4
Understanding	4	5	3	4	5	3	5	4	5	4
Comfort	5	5	4	5	4	3	5	5	5	5
Effectiveness	5	5	4	5	5	4	5	4	5	NA
Enjoyment	5	5	5	4	5	4	5	5	5	5
Overall	4	5	4	4	4	3	5	4	5	4
Preferred System	Buzzy Jacket		Buzzy Jacket		Footboard & Twinkly Lights		Both		Buzzy Jacket	
Preferred Feedback	Vibro		Lights		Vibro		No clear answer		Not directly asked	

Table 7-2: The participants ratings of the two different types of feedback on the smileyometer (5 = happy face, 1 = sad face) and their preferred system and feedback type.

Question	**Alice**	**Bea**	**Carl**	**Daphne**	**Esme**
Prefers playing from memory or notes?	Memory	Memory	Notes	Memory	Notes
Likes to hear a piece played or just plays from written music?	Likes to hear	Likes to hear	Likes to hear	Likes to hear	Likes to hear
Prefers sight reading or aural tests?	Sight reading	Aural test	Aural test	Aural test	Sight-reading
How easy or difficult do you find watching the conductor?*	2	3	3	3	3.5
How easy difficult do you find responding to instructions about posture from teacher?*	5	4	4	5	5

How easy or difficult do you find focusing on particular problem while playing?*	2	4	3.5	3	4
How easy or difficult do you find visualising the ping-pong ball on the fountain to stand up straight?*	4.5	3	3	5	5
What is your approach to correcting a bad habit?	Use visual attention (play from memory)	Think about it, maybe use mirror	Think about it - mainly feel it but might visualise it or use mirror	Think "fingers, fingers"	Say it a lot
How would you design feedback to help straight bowing?	Vibrations + reward/ evaluation at end of piece	Lights with colours having meaning	Vibrations or other tactile	Detailed visualisation on screen with beeps to catch attention	Contributed to Daphne's idea

Table 7-3: Summary of interview responses to questions about the participant's approach to playing music. Questions marked with * used the smileyometer to help the participants express how difficult they found a task (5 = smiley face, 1 = sad face).

Most of the children practised for about 30 minutes about 2-4 times each week depending on their school commitments. Daphne and Esme were more vague about how regularly they practised. Daphne noted *"when I feel like it or when I have the time"* and Esme commented *"whenever I have time".* All the children said they had to be reminded by their parents to do their practice. Participants said that using the prototypes did not interfere much with their normal practice routines. Alice skipped her exercises to focus on playing her pieces with the prototypes and played sitting down when using the footboard. Bea found using the prototypes added a little time to getting ready for her practice.

Four of the five participants had the same teacher and all consistently said that she liked to describe postures or movements and see how the children interpreted what she said. All the children rated this as easy or very easy to respond to. They also refer to her reminding them to correct frequent errors or bad habits by using short phrases such as *"no the fingers!"* while they are playing, in a sense a form of verbal real-time feedback.

Multi-modality preferences

Based on the previous studies it was thought that the students would have strong individual preferences for one modality over another. However, the interviews and ratings showed that unlike the adults in previous studies, the children did not display a strong preference for a particular modality. More important than modality was the learning focus of the prototypes. When asked which prototype was their favourite they chose the one which they felt had the more helpful learning focus irrespective of whether it had their favourite type of feedback. This can be seen in the mismatch between their preferred system and preferred type of feedback shown in the last two rows of the Table 7-2.

One of the research questions this study was designed to answer was whether a person's preference for using a particular modality affected their approach to playing music. Their responses to questions about their approach to music playing (summarised in Table 7-3) showed there is not to be a link between these two. These findings suggest that a multimodal approach may be particularly suited to children as they are not drawn to a single modality but are willing to work with more than one modality. This was also demonstrated in the design section of the interview where two of the children designed systems which used more than one modality.

Motivation

The style of feedback preferred differed between individual participants based on their motivation for playing rather than modality. Participants in this study were motivated by five different factors: competition, performance goals, sense of achievement, aesthetic enjoyment and intellectual curiosity. **Competition refers** to their drive to perform better than their peer group and to engage in games where this judgement could be made. All the children were competitive and found feedback that gave them the opportunity to compete particularly motivating. **Performance goals** refer to concerts or exams where the child has the opportunity to show others their talent. Carl was particularly motivated by performance goals and for this reason he liked the way he felt the vibrotactile forced him to improve his playing which would be beneficial to his next performance. **Playing for a sense of achievement** is more incremental goal where the player enjoys being able to see their progress towards mastering the instrument, for example, through the praise of their teacher or through feedback from technology. Alice was particularly motivated by a sense of achievement preferred the concept of positive feedback or a rating that would allow her to measure her progress. **Aesthetic enjoyment** refers to the direct sensory pleasure that comes from playing and hearing music, as well as that which might come from the sensory feedback in a design. Bea found the lights aesthetically enjoyable to use. This is the most likely explanation for why Bea chose the lights as her preferred form of feedback even though she believed the vibrotactile feedback was more effective for forcing her to improve. **Intellectual curiosity** refers to the pleasure which comes from learning and understanding – Daphne showed particular interest in designing feedback that would let her explore playing from new angles and this reflected her inquisitive approach to the world which she displayed in ensemble practices through her many questions. A full description of the individual differences between participants is given in Appendix E.

7.5 Discussion

This study investigated how multimodal real-time feedback could be integrated as part of a programme of learning about posture with the aim of understanding how real-time feedback can become part of the extended process of learning to play the violin through various forms of practice. By studying real-time feedback in this way, topics of engagement with learning and sources of motivation came to the fore.

The feedback was successful in the way the children engaged with it and with their learning. This can be seen in the way they appropriated the Buzzy Jacket and Twinkly Lights for their own use. When given the Footboard and Twinkly Lights they immediately suggested ways to turn it into games. With the Buzzy Jacket they used the calibration mechanism to find the level of difficulty that was appropriate for them at the time and tweaked this depending on how they were feeling. They also appropriated the video cameras as an audience to perform to. This appropriation shows the children taking control of how they want to learn. Rather than passively following instructions, they actively engaged with making the feedback more appropriate to their practice (tweaking the calibration to make it easier or harder), more enjoyable (turning it into a game) and more motivating (turning it into a performance).

The level of appropriation found in this study was more than that found in the previous study (chapter 6). The reason is that appropriation was dependent on two aspects. First, it was important that the feedback was not stand-alone but part of a programme of learning so the children had the domain knowledge and confidence to take control of their learning. Second, the calibration process enabled them to adjust the feedback and set their own goals. The two elements are now discussed in more detail.

7.5.1 Programme of Learning

The study also shows that real-time feedback can be integrated as part of a larger programme of learning. The initial sessions discussing posture and practicing how to sit and hold the violin were very important to the effectiveness of the feedback and how the students later appropriated it. These initial sessions ensured that before the children were given any technology they had learnt about posture and understood what they were aiming to achieve. Moreover they had already tried other methods such as mental visualisations and exercises to do it. Students needed to know about posture in order to be able to calibrate and interpret the feedback properly. Discussing posture before introducing the feedback made good posture a more significant goal to be working towards than before the start of the study. Understanding posture enabled the students to take control and make judgements about their own learning because they knew what they were aiming for and could appropriate the technology in a knowledgeable way.

The feedback in this study was introduced gradually over time. This approach has advantages. It means technique can be broken down and practised separately preventing students feeling overloaded by having a lot of feedback at once. Once a student is familiar with one feedback system, another can be introduced that builds on it - as in this study. Unlike feedback from a teacher, which is based on the teacher's judgement, with real-time multimodal feedback students themselves can take ownership of the decision about what feedback they have and how they use it. This makes them responsible for their own learning and encourages them to stay engaged. Giving students control of how they get feedback also acts to preserve music playing as largely an intrinsically motivated activity. The students in this study certainly showed enough self-awareness to be able to contribute to this kind of decision-making.

The long term nature of the study meant that when the children started to use the Buzzy Jacket and the Twinkly Lights they were already in an environment where they felt relaxed and comfortable. They knew each other and they knew me and we had built up a relationship where contributing ideas was encouraged. Therefore, they did not feel constricted by the recommended way to use it I gave when I explained how it worked.

7.5.2 Calibration

The calibration processes to individualise the Buzzy Jacket and Footboard and Twinkly Lights were both designed to enable the participants to take control of the feedback and choose how they wanted to learn. The Footboard and Twinkly Lights did not have a regular calibration process. The calibration of the Footboard and Twinkly Lights was carried out through the process of the children drawing around their feet and then having a further two weeks to adjust their foot position and redraw their feet (which several participants did) before the footboards were built. The low-tech approach using paper and pens also had a number of advantages. It allowed the children to set the foot position in a way that was easy for them to understand and directly mapped on to their body. This meant as soon as they saw the new prototypes they understood how they worked and where to place their feet. Through the process of choosing their own foot position they also learnt about what they were aiming for in terms of positioning feet and distributing their weight. Using individual footprints also gave a sense of ownership of the technology. This shows how incorporating familiar materials and mappings (e.g. drawing around body parts) can make calibration very accessible. It may be this familiarity which gave the children the confidence to start suggesting games to play with the Footboard and Twinkly Lights within minutes of being given them.

The Buzzy Jacket had to be calibrated each time it was used. Some students then tweaked the calibration depending on how they were feeling to make the task harder or easier. The calibration process was reliant upon the children having the domain knowledge and the motivation to set appropriate goals for themselves. As already discussed the domain knowledge was provided by the way the feedback was embedded into a programme of learning. In terms of motivation, the students were already motivated to learn their instrument. Theories of intrinsic and extrinsic motivation (deCharms, 1968) suggest that working towards a goal that the child has set themselves will be more motivating than working towards an external goal. This is a strong argument for enabling students to calibrate feedback themselves. Motivation will be discussed further in the next section.

Another aspect of the calibration of the Buzzy Jacket that was important was the way that it was a quick process which took only seconds to achieve. Real-time vibrotactile feedback from the Buzzy Jacket afterwards meant that it was easy for the students to quickly judge whether the calibration was comfortable or not. This meant that the children could tweak their calibration as much as they wanted without having to invest much time. It also meant they could quickly learn how their action in the calibration process related to the outcome for the feedback afterwards. In the terminology of Dix's (2007) guidelines for appropriation the fast calibration process combined with the real-time feedback afterwards *"provides visibility".*

7.5.3 Motivation

It is common to classify motivation as either intrinsic or extrinsic and certainly we can see that practicing for a concert is a more external source of motivation than practicing purely for aesthetic enjoyment or intellectual interest. This classification is useful because studies show that intrinsic motivation tends to lead to more sustained engagement and is generally considered preferable to extrinsic motivation (e.g. (deCharms, 1968), (Ryan & Deci, 2000)). In some cases the addition of an extrinsic motivator can make a task less intrinsically enjoyable (e.g. (Calder & Staw, 1975), (Greene & Lepper, 1974)). However, extrinsic motivators which represent achievement can increase intrinsic motivation (Hallam 2002), (Ryan & Deci, 2000).

However, using de Charm's definition of intrinsic motivation (de Charm 1968) which is when *"a person experiences himself to be the locus of causality for his own behaviour"* as opposed to extrinsic motivation "when a person perceives the locus of causality for his behaviour to be external to himself" the distinction between motivators is not clear cut. For example, Carl who was motivated to practise by upcoming performances, deliberately sought out opportunities to perform because he enjoys performance. Therefore while the performance is an external source of motivation, Carl is still the locus of causality for his behaviour. Similarly, Alice's motivation to practise in order to have her achievement externally recognised through praise or a good rating was an expression of self-efficacy rather than her being controlled by her environment. Moreover, extrinsic motivators which signify achievement generally do not have a negative effect on intrinsic motivation and in some cases can increase motivation (Hallam, 2002).

Competition is more clearly extrinsically motivating because an individual is striving to do better than their peer group and thus the individual's actions are controlled by the actions of others. On the other hand, many of the children were proactive in creating situations where they could enter into competition with one another.

One case where making a distinction between extrinsic and intrinsic feedback is useful is when considering the concept of *forcefulness* of feedback. Two of the participants (Bea and Carl) described the vibrotactile feedback as forcing them to change their position. Feedback that gives a sensation of being made to do something clearly fits the definition of an extrinsic motivator. Extrinsic motivators like this can potentially damage intrinsic motivation to practise. However, to discount vibrotactile feedback on these grounds would be an over simplification. This is where the student's own motivations to practise become important. How the feedback affects motivation is dependent upon the child's reasons for using the feedback. For example Carl valued the forcefulness of the feedback because it helped him in his practice for performances which he found intrinsically motivating. Whereas Bea would not choose to use vibrotactile feedback in her future designs because she practised for different reasons and wanted to find the practice itself enjoyable. If students have made a free choice to use forceful feedback and can control its strictness, then they remain the locus of causality and are just using the feedback as a tool in something they are intrinsically motivated to do. If students feel they ought to use the feedback, then the forcefulness of the vibrotactile feedback could become problematic. Therefore, it is not just the nature of the feedback itself that is important it is the setting it is being used in and the student's personal motivation.

A similar argument can be made about using competition to motivate students. If competition is freely entered into for enjoyment like the games suggested by the children in this study then the locus of causality remains with the student. Therefore, there is no reason not to utilise the motivating power of competition. However, if students feel under pressure to compete then they will feel controlled by their environment making them extrinsically motivated. This means that when designing feedback, competition should not be an integral part of the design because that would force students to compete. However, a design which has the option to facilitate friendly competition could utilise the motivating power of competition without damaging intrinsic motivation.

Another distinction that is made within education is the difference between performance related goals and mastery related goals ((Ames & Archer, 1988), (Elliott & Dweck, 1988)). Performance related goals refer to getting positive recognition of competence and avoiding being viewed as incompetent. Mastery or learning goals on the other hand focus on learning a skill and gaining mastery rather than having that mastery recognised. Mastery goals have been shown to lead to better learning strategies ((Ames & Archer, 1988), (Smith, 2005)) whereas performance goals can lead to a helpless response in some cases (Elliott & Dweck, 1988). In this study the children were motivated by performing well in front of others and having their achievement recognised this can be seen in their competitive interaction with one another. However, they were also motivated by progressing their learning. For example, Alice was keen to have praise from teachers which she valued as a measure of her learning. Her ideas of a rating scale would be another way which she could measure her improvement. She was motivated by achieving goals and feedback in the form of a rating presents the opportunity to encourage her to focus on mastery goals rather than performance goals.

Within music learning we must also be cautious of condemning performance goals because public performance is an integral part of the subject. Carl enjoyed performing for others and this experience was not only acted as a goal to achieve but also an experience to relish. So while performance goals can be counter-productive, most music teachers would agree that working towards performances is part of musical training and the experience of performing can be inspiring (Sloboda, 2004).

Competition should also be viewed as a form of performance related goal as it concerned with beating others on whatever criteria rather than mastering a skill. As such one might expect it to be demotivating if an individual was not well matched to their peer group. Additionally, a learner may not adopt the most appropriate learning strategies if they are too focused on competing with others.

In sum, all these factors (competition, performance goals, achievement goals, aesthetic enjoyment and intellectual interest) are powerful motivators to practise. Categorising them either in terms of intrinsic and extrinsic motivation has raised important design issues when considering using forceful feedback or competition. These issues can be overcome by leaving enough flexibility in the design so that the student can control how they use the feedback. Categorising in terms of performance and learning goals also raises the design issues concerning competition and achievement rating. Concerning ratings it is important when giving rating that they are focused on learning goals and used as a way of tracking learning rather than as an aim in themselves. Competition can also lead to performance related goal and needs to be confined to situations where peers are equally matched and must remain only a minor part of the design.

Rather than designing feedback to promote a single type of motivation, a programme of multimodal feedback has the opportunity to enhance the many motivations children already have to practise. Feedback can be designed to enhance aesthetic enjoyment by adding pleasant sensory feedback such as lights or visual imagery or potentially even music to add to the experience of playing. Feedback can be intellectually engaging by offering players a novel perspective on their playing. In this study the feedback acted to make players become aware of their posture and how they changed it while playing and learn things about their own movement which they had not known before. A more complex display has the potential to be more intellectually engaging – for example the screen display Daphne suggested. However, there is a trade-off between making something intellectually engaging and making something which can be used peripherally while focused on playing music.

7.5.4 Extending the Multimodal Real-time Feedback Framework

Chapter 5 presented a model of real-time feedback around the processes of catching attention, interpreting the feedback and taking action. In the interviews in this study these processes also showed where the technology interventions fitted the framework. Lights could be difficult to notice while playing music for some individuals whereas others reported having difficulty feeling the vibrotactile feedback. The lights were more difficult to interpret than the vibrotactile feedback (so long as player could perceive the location of the vibration) for two reasons: the vibrations are located on the body drawing attention to the particular area on the body that needed to respond; the lights only pointed out a mistake and gave no additional information on how to correct foot position.

Preferences towards different modalities were less extreme than those expressed by adults, indicating that children may be well suited to being given feedback in multiple modalities. Children generally valued the relevance of the learning goal over the modality of the feedback when choosing their favourite prototype. Other characteristics of the feedback such as whether they gave positive or negative feedback and whether they were enjoyable to use were also very important to the children. The particular characteristics that were valued were related to the children's personality and their personal motivations for playing the violin or viola.

Bea and Carl both described a characteristic of real-time feedback that has not been described so far in the model of real-time feedback. This could be termed the *forcefulness* of feedback. This relates to the taking action stage of the framework. Both children describe how they felt the vibrotactile feedback was more demanding for a response than the visual feedback making it in their view more effective at changing their playing. Even though Carl described the vibrotactile feedback this way he also described sometimes not being able to feel the feedback as well. This shows that it is not just about how good the feedback is at catching attention, forcefulness is a different quality connected to how much it galvanises action. The vibrotactile feedback stimulates action because it is a sensation that the player wants to respond to; whereas the lights can be deliberately ignored. Although the concept of forceful feedback may initially appear desirable it can have drawbacks too as it has the potential to effect student motivation.

7.6 Conclusions

How real-time feedback influences motivation appears to be an interaction between the individual's personal sources of motivation, the demands of the setting and the characteristics of the feedback. This means that feedback cannot be directly designed to instil motivation in players but it can be designed with flexibility to enable users to draw on their own intrinsic motivators. All these factors are discussed further in the next chapter which describes more fully a theoretical framework to explain how real-time visual and vibrotactile feedback can be designed to facilitate and sustain the learning and practice of sensory motor skills.

7.7 A Final Word on Methodology

The long-term in-the-wild nature of this study with close researcher involvement has enabled the investigation of children engaging with feedback technology in the real-world. In this study technology had to contend with other day-to-day concerns such as homework, school clubs, broken violins and television just as normal violin practice has to. Many of the findings were unexpected. For example: the use of the video diary for performance; the competition among the group and how this fed into how they wanted to use the feedback; and the appropriation of the calibration method for setting difficulty. Although this was challenging in terms of data collection, this is the also the value of an in-the-wild study because it reveals what the important and relevant issues are in the real-world and gives findings that are relevant for design.

In the interviews, the use of the design exercise question proved particularly revealing because it provided an opportunity for the children to express their ideas about technology for use in practice without the constraints of referring to the designs used in this study. This style of questioning *"to discover what we did not know rather than try to confirm what we thought we already knew"* (Scaife & Rogers, 1998) is fitting for the philosophy of in-the-wild studies which explore the unknown and unanticipated.

My role in the research was very hands-on and similar to that of a teacher when running the ensemble practice sessions. In this position I could champion the need for good posture and ensure that participants had all the knowledge required for using the technology. The role of teacher also enabled me to build relationship with the student as learners and discover more about their approach to playing and motivations as well as their strengths and weaknesses in terms playing. This both informed the design of the feedback (e.g. the Footboard and Lights) and the interpretation of the findings.

Chapter 8 – Framework to Describe Visual and Vibrotactile Real-time Feedback for Motor Learning

8.1 Introduction

In his notion of calm computing Weiser (Weiser & Brown, 1996) makes a distinction between the periphery and the centre of attention and discusses how having information available on the periphery which can be centred in on by the user at time of need can be empowering and calming. However, in high demand situations such as playing a musical instrument how do we provide feedback on the periphery that enables learners to improve, rather than leaving them unaware of the feedback or overloaded by it?

This chapter aims to address this question by drawing together the findings from the research conducted in this thesis and drawing on existing models of interaction to structure them into a framework for understanding and designing real-time feedback for sensory motor learning. It builds links between the individual findings of each study to form a more generalizable theoretical framework which can be applied to the future studies of real-time feedback. It provides a new contribution to our knowledge and understanding of real-time feedback which is applicable to many applications such as learning sports, music and for physical therapy.

This chapter is organised into two parts that reflect the two elements of the framework: the first focuses on the stages involved when learners use real-time feedback; the second focuses on the many different factors that influence how these stages take place. Norman (Norman, 1986) describes how when interacting with technology users go through seven stages of action which takes the interaction from a goal in the user's mind to an outcome performed by the technology and perceived and interpreted by the user. In this process two gulfs must be overcome, the gulf of execution and the gulf of evaluation (Figure 3-1). However, as discussed in the literature review, real-time feedback is different because the technology initiates the interaction, therefore the gulf of evaluation comes first. This requires that the feedback must catch attention and then be perceived and interpreted by the user before the user then takes action. The action that the user takes aims to use the information from the feedback to achieve learning goals rather than to interact with the technology as in the action steps in Norman's gulf of execution. As such the stages in using real-time feedback are different to those of Norman but are informed by them. The framework of real-time feedback for motor learning has four stages: catching attention, taking in information, interpreting it and acting (Figure 5-34).

Although it is useful to think of these as separate stages of interacting with real-time feedback, the Theory of Event Coding (Hommel, Müsseler, Aschersleben, & Prinz, 2001) suggests that observed, intended and executed action are all encoded in the same way and can facilitate and interfere with one another. The research in this thesis supports this hypothesis. For example, it was found that the intention to take action and change playing using the feedback helps players to be more aware of the feedback and so it catches their attention better. Therefore, in the framework of real-time feedback for motor learning, the stages can be thought of as interlinked so that different stages influence one another.

Many factors influence how the stages shown in Figure 5-34 take place and whether they flow from one to another easily. The second part of the framework of real-time feedback for motor learning models these factors (Figure 8-1). Key to the use of real-time feedback is the learner's embodied presence within a physical and social context. Embodied interaction (Dourish, 2001) places strong emphasis on the way that interaction with technology is not simply controlled by its design, but the way users interact within the world. This is strongly supported by the research reported in this thesis which has found that both the social and learning context play an important role in the way feedback is used as well as individual differences is the way users approach the world. Moreover, the proposed model of the factors affecting feedback use shows how design (feedback characteristics), context (external factors) and individual differences interact to influence how real-time feedback is used and experienced.

In the next sections, first the four stages in part one are summarised and then factors in part two of the framework are explained in more detail, drawing on evidence from the studies in this thesis as well as studies of the MusicJacket and other real-time feedback in the literature as examples. In order to enhance clarity Table 8-1 gives a summary of the topics covered in each of the study chapters.

Chapter	Setting	Topics covered
Initial Work (van der Linden, Johnson, Bird, Rogers, & Schoonderwaldt, 2011)	Violin Lessons with Children	How real-time vibrotactile feedback can be incorporated into violin lessons
5	Individual practice of orchestra players MuSense	Comparison between visual and vibrotactile feedback Individual differences in response to these modalities
6	Ensemble practice at a summer school ShareSense	Real-time feedback in ensemble playing Shared visual feedback
7	Ensemble and home practice of children Buzzy Jacket and Twinkly Lights	Real-time feedback as part of a programme of learning Individual differences in motivation

Table 8-1: Summary of the empirical studies of real-time feedback in this thesis

8.2 Part 1: Stages in Using Real-time Feedback

As identified in Chapter 5 there are four processes that take place when participants use real-time feedback (Figure 5-34). These are *catching attention, taking in information, interpreting* and *taking action*. These are summarised below, a more detailed description of these factors can be found in 5.6.

8.2.1 Catching Attention

In order for real-time feedback to be used it must catch the attention of the player while they are playing so that they are able to consciously become aware of something that is communicating with them. Whether real-time feedback catches the attention of the player is influenced by many factors: the *conspicuousness* and *positioning* of the feedback and the individual's *peripheral awareness* while they are playing; how demanding the player is finding the playing task; whether the *feedback focus* fits the *learning objectives* of the player; whether the *modality* of the feedback matches an individual's ability to perceive it; and whether feedback is *publicly visible* in a *social context* where others might act to point out the feedback.

8.2.2 Taking in Information

Once a player is aware of the feedback they need to be able to take in what it is showing in relation to their own playing. In the case of visual feedback they need to see the whole display and also be aware of their own movement which caused it. In the case of vibrotactile feedback they need to be able to identify what motor was triggered, when and what movement they were making at the time. The ability to take in information is affected by: the *complexity* and *position* of the display and an individual's *peripheral awareness*; how difficult the player is finding the task; and the *modality* of the feedback and how this relates to the player's *modality preferences*.

8.2.3 Interpreting

A player also needs to be able to interpret this information to understand what it means for their own playing. For example, when using the Buzzy Jacket, players need to understand that the vibration they are feeling on the base of their elbow means that they have let their violin drop down. The way people interpret the feedback depends on: the *metaphor* and *modality* the feedback uses to communicate information and how this relates to their individual *approach to playing* and *modality preferences*; it is also influenced by the *learning context* and how much they understand the objectives.

8.2.4 Taking Action

Based on the interpretation of the feedback the player must then take action to adjust his or her playing. For example, with the Buzzy Jacket, if a player feels a buzz on their elbow they need to lift their violin up. To act upon the feedback, players need to know what action to take and also be motivated by the feedback. Knowing what action to take is dependent upon the interpretation and also influenced by the *learning context* in terms of understanding what movement they are aiming for.

Participants must also be motivated to take action in response to the feedback. Motivation is influenced by the relationship between the *focus of the feedback* and their personal *learning objectives* and the *learning context*; and the *social context*, particularly if the feedback is *publicly visible*. Player's *individual motivators* and the way that they interact with the characteristics of the feedback also play a role in motivation. For example, players that are motivated by a sense of achievement are more likely to be motivated by positive feedback that allows them to see their improvement through the feedback.

8.3 Part 2: Factors Influencing How Real-time Feedback is Used

Many factors influence how real-time feedback is used as shown in Figure 8-1. These factors were identified by examining the findings from the three studies in this thesis and the previous work with the MusicJacket (van der Linden, Johnson, Bird, Rogers, & Schoonderwaldt, 2011) and comparing the way that different designs of feedback functioned in different settings and with different people. They are organised into three categories: Feedback Characteristics which concern how the feedback is designed; External Factors come from the context the feedback is used in; and Individual Differences which are dependent on the player using the feedback. The factors interact with one another. For example feedback focus – e.g. long bowing (feedback characteristic), task – e.g. whether the music needs long bows (external factor) and personal learning objective – e.g. whether the player wants to learn to use long bows (individual differences) will all interact to govern whether the feedback is relevant to the player carrying out a particular task and this in turn will influence how motivated they are to take action on the feedback and what resources they dedicate to monitoring the feedback. These categories are useful because they separate out what is in the decision of the researcher/designer and what needs to be accommodated in the design. The characteristics of the feedback are in the designer's choice; the designer chooses which modality to use and how to present information. Whereas, external factors like social context and individual differences like modality preference are not in the designer's decision space and ideally need to be accommodated in the design.

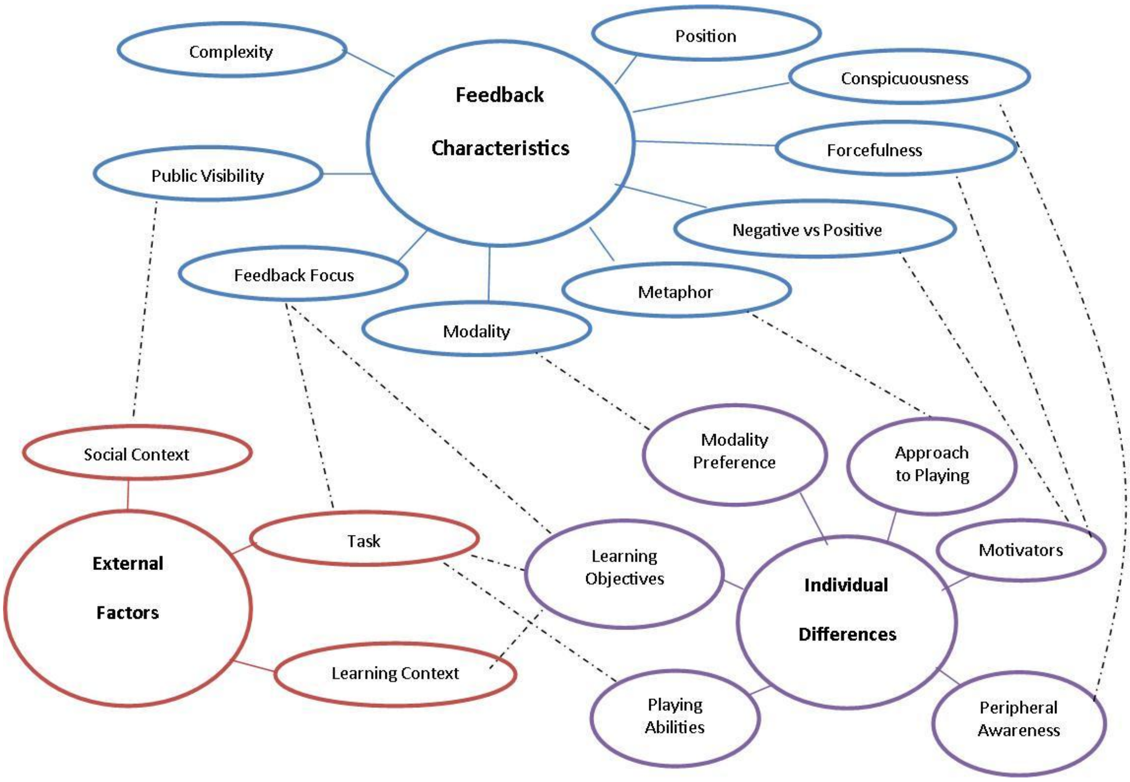

Figure 8-1: Factors influencing how real-time feedback is used with some of the most common interactions between factors (dotted lines).

In different situations different factors in this framework come to the fore. The different studies in this thesis revolved around different key factors depending on the study design and the context.

8.3.1 Feedback Characteristics

Certain characteristics of the feedback can have a large influence over its use. The design space for real-time feedback devices is huge and there are a myriad of individual designs all of which will have different pros and cons. Therefore, I have divided up this space in terms of key characteristics of the feedback which the research in this thesis has shown to be influential in the way feedback is used. These characteristics are summarised in Table 8-2. Each characteristic will now be considered in turn.

Feedback Characteristic	Description
Feedback Focus	The aspect of playing technique that the feedback is being given about.
Modality	The sense being used to give feedback through – e.g. tactile, visual, auditory.
Conspicuousness	How much the feedback attracts involuntary attention.
Positioning	Where the feedback is positioned. For vibrotactile feedback this will be places on the body, for visual feedback this may be anywhere in the environment.
Metaphor	The analogies used to help people understand and remember the feedback. For example being 'pushed' or 'pulled'.
Positive versus Negative Feedback	Whether the feedback points out goals achieved (positive feedback) or mistakes that are made (negative feedback)
Public Visibility	Whether other people can also see the feedback, or whether it can only be perceived by the player.
Complexity	How much information the feedback aims to communicate.
Forcefulness	How much the feedback is perceived as forcing a response from the player or making them improve.

Table 8-2: Summary of feedback characteristics

Feedback Focus

Real-time feedback for learning a physical skill needs to have a well-researched aspect of technique to focus on. Early studies with the MusicJacket showed that having feedback about several aspects of technique in one modality was difficult for users to respond to as it is difficult to take in where the vibrations were coming from and participants felt overloaded (Johnson, van der Linden, & Rogers, 2010). Therefore it is best to keep feedback focused to a single element of technique at a time, or to use two modalities (cf. Chapter 7) to give feedback about two separate elements of playing.

All the studies show that for feedback to be useful it must be relevant to individual learning objectives or fit into the learning context – e.g. fitting with the objectives set by an expert in the eyes of the learner such as a teacher or conductor (e.g. Chapters 5 and 7). This provides motivation to use the feedback because the outcome of using the feedback has value to the player. When the feedback is not relevant to the learner's objectives, learners are less likely to act on the feedback. The study at the high school reported in chapter 7 showed that feedback which facilitated learners to set their own goals within a particular feedback focus was particularly engaging.

The relationship between the feedback focus and the task that the player is working on is also important. If, for example, the player is focused on learning the notes but is being given feedback about long bowing as in the second study then this makes using the feedback difficult. The player has to think about the task focus and the feedback focus at the same time and can make it difficult for feedback to attract attention.

Modality

Different modalities have different properties and are different in the way people perceive and understand them. The first study showed that it is not the case that one modality is superior to another for giving real-time feedback. Overall in terms of the effectiveness, vibrotactile and visual feedback were found to be equal. However, on an individual level, different people respond differently to different modalities and these differences are substantial. What goes into creating these individual differences will be discussed in more detail in the Individual Differences section.

Given the individualised responses to different modalities it is hard to make generalisations about the characteristics of particular modalities, but there are some patterns. People who liked the vibrations best felt that they gave a sense of “knowing” which appears to combine the processes of catching attention, taking in stimulus and interpretation into a single experience. Whereas those who liked the visual feedback give more detailed analytic descriptions of the process of interpretation such as relating it to their own mental visualisation of the bow or relating it to the written music or a sense of achievement from a ‘good’ stimulus. However, it is important to note that people are not always able to tell which modality is having the largest effect on their playing. Some of those reporting this particular affinity for vibrotactile feedback did not get the best objective results when using it. It appears that the blurring of boundaries between catching attention, taking in information and interpreting may reduce the amount of conscious attention which is paid to using the feedback because less effort is required to interpret it; this results in a better perceived experience but not necessarily the most effective feedback.

Initial ideas when designing the feedback and selecting a modality suggested that visual feedback may be unsuitable as it may interfere with reading the music. Although this did happen to an extent, we found that cognitive load seemed to play a larger role, and this affects all modalities. For example, in the ensemble practice at the summer school in chapter 6 participants had difficulty being aware of either modality, as ensemble playing is a particularly demanding situation. Whereas in the study of individual practice in chapter 5, which compared visual and vibrotactile feedback, participants did not have difficulty responding to either visual or vibrotactile feedback while reading music as they were in a less cognitively demanding situation.

Conspicuousness

Feedback needs to be noticed to be useful; the first step in my model of real-time feedback is *catching attention*. However, feedback should not be so conspicuous that it prevents participants from concentrating on other parts of playing. The key here is to mediate attention in an appropriate way. Certain characteristics of feedback act to catch attention. In particular, discrete changes in a display such as flashing on a visual display or turning on a vibration motor are more noticeable than continuous changes such as gradually building up lights or adding a second vibration motor. These are known as exogenous cues (Mulckhuyse & Theeuwes, 2010). If a designer wants to draw attention to the feedback, then sudden changes are necessary. However, there is a trade-off between noticeability and the complexity of information being communicated, as more information can be encoded in continuous changes than discrete ones.

Conspicuousness is also dependent upon positioning. A visual display positioned close to the main visual focus, in this case the sheet music, is much more noticeable than a display projected onto a wall. Vibrators positioned on certain parts of the body are more noticeable than on other parts of the body. It is also important to note that there are individual differences in terms of people's peripheral awareness so there is no single ideal level of conspicuousness as this varies with the individual and the setting.

Two studies were conducted which involved giving feedback in two modalities at once. In the summer school study reported in chapter 6, which studied ensemble playing, participants were given visual and vibrotactile feedback about the same element of playing - bowing. This did not appear to make the information more noticeable than using a single modality. In contrast, in the third study two modalities were used to give feedback about two elements of playing. Spreading the information across two modalities meant that feedback could be designed in a particularly noticeable way because less information had to be encoded in each modality. Theories which postulate that there are dedicated sections of working memory for each perceptual channel ((Baddeley, 1992) (Mayer & Moreno, 1998)) would argue that this multimodal approach would also be easier to notice, take in and interpret.

Positioning

The position of the feedback is important not only for catching attention but also for all stages of the real-time feedback model including suggesting action. The position of vibrotactile feedback changes how easy it is to notice and then perceive its location; certain positions on the body can also suggest certain actions. This is dependent upon nearness to reference points on the body such as joints and the localised resolution of skin sensation and the underlying tissue which the vibrator is positioned against ((Cholewiak, Brill, & Schwab, 2004), (Cholewiak & Collins, 2003), (Brown, 2007)). The main approach used was to place vibrators on bones rather than muscle for comfort and position them close to joints for clarity. This worked well with all participants learning how to respond to the vibrators and some participants commenting on how the positioning of the vibrations on the body helped them to understand how they needed to change their movements. This shows that the position of vibrators on the body can be used to communicate movement.

The position of visual feedback is also very important. Feedback needs to be positioned in the peripheral vision in order to be able to attract attention (Lamble, Laakso, & Summala, 1999). If it is positioned too far away, as in the summer school study in chapter 6, it cannot mediate attention in the same way. On the other hand, if visual feedback is designed to be shared by many people, or be publicly visible in order to encourage peer motivation then a trade-off needs to be made in terms of position. One successful solution is the design in the study with high school students in Chapter 7 where the fairy-lights were on the music stand but also visible to those around. This was successful in catching the attention of individual learners and their peers, meaning that the feedback could be used socially as well as personally.

Metaphor

Feedback can use metaphors to help participants to learn how to interpret it. The ease of learning to respond to the feedback is dependent upon the clarity of these metaphors and the way they link to the metaphors individuals already use when playing the violin. The designs in this thesis have used a range of metaphors to help explain the feedback. The earlier work with the MusicJacket (van der Linden, Johnson, Bird, Rogers, & Schoonderwaldt, 2011) and the Buzzy Jacket both used the *push* metaphor, where participants were told that if they felt a vibration they needed to move away from it. Even with this metaphor, it still takes time participants to learn the correct response to the feedback because there are many degrees of freedom in the body. Some of the school children using the MusicJacket (van der Linden, Johnson, Bird, Rogers, & Schoonderwaldt, 2011) moved their elbow up and out to the side when they felt the vibration on their left elbow (violin side); this fits with the push metaphor, but was not the desired movement which was to lift up their whole arm. This shows that while a metaphor can be useful, it is not a complete explanation in itself. This was overcome in the Buzzy Jacket study in Chapter 7 by making the feedback part of a programme of learning about posture so that the children learnt the movements they needed to make before using the feedback.

Another type of metaphor that participants used to think about the vibrotactile feedback was of good or bad *zones* which they moved their body into. For example, participants using the MusicJacket or the Buzzy Jacket could also view the feedback about violin hold as vibrating when their elbow moved down into a *bad zone* and switching off when they moved their elbow up and out of the *bad zone*. Similarly, with the long bowing prototypes reported in Chapters 5 and 6, participants could imagine the vibrations as coming on when their hand moved into a certain *good zone* in space which is the place they reach at the end of a long bow. This idea of imagining the vibrations coming from the environment rather than being on the body is interesting because it may encourage an external locus of focus which can improve performance (Weiss, Reber, & Owen, 2008). Further research would be necessary to find out whether the *push* metaphor or the *zones* metaphor change people's locus of focus and whether this can improve performance.

Metaphors were important in the visual feedback. The complex visual feedback in chapter 5 showed the bow length building up as a line. One participant said they found this particularly easy to understand as he already visualised the bow as a line. Another referred to it as a *"timeline"* which she found easy to understand. This shows that learners already use metaphors for thinking about playing and that feedback can capitalise on these to make the display easier to understand.

The studies also show that people can learn new metaphors quickly. The visual feedback in Chapter 6 used a flower as metaphor for the group and each petal as the playing of an individual. This is an unusual metaphor, none of the participants indicated that they had thought of their playing in this way before, but still all participants learned to understand quickly that their bow length related to the length of the petals. The real-time continuous response to their movement aids learning, also the dimensions are consistent in this mapping - a length corresponds to a length. Whereas not all the participants understood the way the speed of the bow changed the width of the petals which is a case where the dimensions are not kept intact in the mapping. When selecting a metaphor for visual feedback it may be clearest if dimensions are consistent where possible.

Positive versus Negative Feedback

In addition to metaphors feedback can also be used in a positive way to indicate 'good' playing or a negative way to indicate 'bad' playing. The original MusicJacket (van der Linden, Johnson, Bird, Rogers, & Schoonderwaldt, 2011) and the Buzzy jacket both used negative feedback, which told the learner when they were making a mistake. Having feedback which told participants when they were doing something wrong may be disheartening as it encourages learners to dwell on their failures rather than their progress, which could for example give them low expectancy of success in an expectancy x value model of motivation (Eccles & Wigfield, 2002). However, it was found that with the Buzzy Jacket and the Twinkly Lights, which used negative feedback, learners interpreted no feedback as success, and showed a sense of achievement. Therefore, it is too simplistic to only look at whether the feedback stimulus is triggered by doing something right or wrong as learners also treat the absence of feedback as a result. So long as learners are able to measure their progress using the feedback, they can use it to motivate themselves. The designs reported in the first and second studies used positive feedback to indicate when the player was doing something correctly. Interestingly, when participants were interviewed in the first study about this aspect of the feedback, some participants said that the green lights in the visual feedback gave a clearer sense of doing something "right", whereas the vibrotactile feedback felt less like a value judgement. Many participants also proposed negative feedback for future designs and did not feel that this would affect their motivation.

There are individual differences in attitudes about positive and negative feedback. One child in the high school study in Chapter 7 preferred feedback that praised achievement rather than pointing out mistakes. She was still happy to use vibrotactile feedback as negative feedback but would have also liked a reward or a rating to measure of her progress in order to balance it. This shows that while some people are comfortable with negative feedback a more positive approach might be more suitable for others. This difference likely comes from individual differences in participant's goals and the intrinsic and extrinsic motivators that encourage them to practise. These are discussed in more detail in the Individual Differences section.

Public Visibility

Feedback can either be personal, only visible to the individual learner, or it can be public, visible to learners, their peers and their teacher. Vibrotactile feedback is inherently personal whereas visual feedback can be publically visible. Previous studies with the MusicJacket (van der Linden, Johnson, Bird, Rogers, & Schoonderwaldt, 2011) found that vibrotactile encouraged discussion between students and their teacher because the teacher had to ask the pupil what they had felt, which meant that the pupil started analysing their own movement. This finding was confirmed in the high school study with the Buzzy Jacket where the teaching relationship change after vibrotactile feedback was introduced. This is an important consequence of having feedback which can only be felt by the learner.

In contrast, visible feedback with groups of players was tested publically. The success of this approach was dependent upon the social context. Most notably, it was observed that among children who were already competitive with one another, the public visibility of the feedback fed this motivating force and they tried to get better feedback than the others. In terms of an expectancy x value model of motivation the potential of the feedback for publically showing achievement increases the value of the achievement. There was a concern that having progress shown publically could make children anxious so we asked the children whether they minded having the feedback visible to others and all responded that they did not. This may be because playing is already a very public activity, they cannot play in the small group they were practicing in without being heard by others, and the feedback is simply an extension of this. However, it may be that in a different social context, participants may feel less comfortable with public feedback. Personal and public feedback both have advantages and which is most suitable is dependent upon the social context which they will be used in.

Complexity

The amount of information that real-time feedback is trying to communicate will affect how well participants are able to take it in. While a good metaphor may make large amounts of information easy to interpret, there still may be difficulty perceptually in taking in the information peripherally. For example, some of the university orchestra players using MuSense found it difficult to take in the complex lights which built up and changed colour because they found trying to take in a length and colour was difficult while looking at the music whereas simply taking in whether the simple lights were switched on or not was easier. Similarly, as described in Section 3.7 early studies with the MusicJacket found that having too many vibrators on was difficult for participants to take in which vibrators were active at one time. However, there are also advantages to displaying extra information, it makes a display useful to more people because it gives more ways it can be used. There is a trade-off which has to be made between the amount of information that real-time feedback displays and people's ability to use it peripherally.

Forcefulness

Forcefulness of feedback is how strongly the feedback is perceived as insisting that the player takes action in response to the feedback. It was found in the high school study in Chapter 7 that the vibrotactile feedback was perceived as being forceful whereas the visual feedback was not. This does not necessarily mean that forcefulness is only a property of vibrotactile feedback, it may be that visual feedback can also be designed for this. It is only a property of negative feedback set-up as if feedback is only pointing out correct actions it would not be present to force someone to change their action when they make a mistake.

Whether forcefulness is a good or bad quality is dependent upon the individual. Participants who are motivated by performance goals may value forcefulness because it gives a sense of being made to improve which will feel reassuring when participants are concerned about exams or upcoming performances. On the other hand, those who are motivated mainly by intrinsic enjoyment may find the sensation of having their choice and self-expression taken away by forceful feedback demotivating. In either scenario it is important that participants feel that they have freely chosen to use the feedback and that they are the locus of control for this decision otherwise the sense of being controlled by the feedback may damage motivation (Calder & Staw, 1975).

Feedback Characteristics – Summary

The research in this thesis has demonstrated that there are many characteristics of real-time feedback, giving a large number of options to a designer. Many of these decisions involve making trade-offs between characteristics, for example between complexity and conspicuousness or positioning and public visibility. However, these are not the only factors that impact on the way real-time feedback is used. Both external factors and individual differences will influence how these feedback characteristics act to catch player's attention, allow them to take in information and interpret it and then whether they take action. The next section discusses the external factors which come from the environment that the feedback is being used in.

8.3.2 External Factors

The way feedback is used is influenced by a number of factors which are external to the design itself and the individual using it. The key external factors encountered in this research are summarised in Table 8-3. Although these factors are not directly in the designer's control, feedback can be designed to take these factors in to account. A designer could choose to design feedback for a certain task or group of tasks; similarly feedback could be designed to sit within a particular learning programme or teaching method.

External Factor	Description
Task	The activity that the player is carrying out while using the feedback, e.g. playing a familiar piece of music or sight-reading.
Social Context	Whether there are other people there. The opportunities and demands that come from practising with others.
Learning Context	What the player is currently learning and how they are learning it.

Table 8-3: Summary of external factors that influence how feedback is used

Task

In all the studies reported in this thesis the task participants were working on has played an important role in shaping how feedback was used. The tasks participants had to carry out came from different sources depending on the study. For example, in the last study the tasks came from the concert repertoire of the high school orchestra the children were members of. It is important when conducting in-the-wild studies that the task is valuable and meaningful to participants, otherwise there is no motivation for them to work on it. Task affects the usability of the feedback in several ways.

First there is the focus of the task and how this relates to the focus of the feedback. For example in the study with orchestra players in Chapter 5 the task as set by the conductor was to use full bows in a particular piece and this was directly related to the focus of the feedback which made the feedback fit into the practice process more easily. Whereas the task in summer school study in chapter 6 was concerned with learning to play a piece together and feedback about bowing was not central to this. When the task is not directly related to the focus of the feedback then the feedback will have a more peripheral role and it may be hard for players to pay it attention.

Second the difficulty of a task affects how much attention participants can give feedback. For example, in the summer school study playing in an ensemble was particularly demanding and this may account for the reports of not noticing both types of feedback. This was also found in studies with the MusicJacket (van der Linden, Johnson, Bird, Rogers, & Schoonderwaldt, 2011). The difficulty of a task is specific to individuals and dependent upon their playing experience and how long they have been playing the piece.

Social Context

Social context can help to catch participants' attention and motivate them to take action. The study in Chapter 7 shows the children coming up with games and competitions with the visible feedback which would motivate them to work harder on their playing in order to be the best and win. These children were already motivated by competing with one another and the feedback supported this as a measurement of playing skill that was visible to players. The children in the group also used the visual feedback to remind one another when they were making a mistake, so that if the feedback hadn't succeeded in catching the individual student's attention, then their neighbour would point it out to them. This adds an additional route by which the feedback can catch the player's attention.

In contrast, some social contexts can make tasks more difficult. For example in the study in the summer school study in Chapter 6, playing in the ensembles at the summer school added extra demands such as listening and keeping time with the others which made using real-time feedback very difficult for the participants. In addition to the cognitive demands of ensemble playing, the stress of playing and being heard by others may also impact upon this, for example studies show that stress limits peripheral vision (Williams & Andersen, 1997).

Learning Context

Learning context is also very important to how feedback is used. If students have been gradually learning a particular technique over several weeks, then feedback which is related to this will have more meaning for them than feedback which is tangential to their current topic of learning.

For example, the feedback in the high school study in Chapter 7 was particularly well accepted by students because it was part of a larger programme of learning about posture. Having the feedback in this context was advantageous because the children knew how to interpret the feedback and what actions to take when they felt the feedback because they had already learnt what they were aiming for before using the feedback. Compare this to the summer school study in Chapter 6 where the feedback was given to participants without knowledge of the learning context which made it much harder for participants to decide how to use the feedback. Learning context also plays a large role in influencing individual learning objectives can affect both awareness of feedback and motivation to take action to improve based on the feedback.

External Factors – Summary

The in-wild-studies in this thesis have shown that external factors that come from the context and the task have a large impact on the way real-time feedback is used. This makes different feedback characteristics more suitable for different types of situations. It also means that real-time feedback can facilitate learning in different ways depending on what is appropriate to the context. However, external factors are not the only influence on how real-time feedback is used, equally important are the characteristics of the individual using it.

8.3.3 Individual Differences

On an individual level people use and experience real-time feedback very differently. The modality comparison study in Chapter 5 found that although there was no difference between three different designs of real-time feedback when averaged over all participants, there were large variations between the effectiveness of different designs for each individual participant. This section discusses the many differences between people that can affect how they use real-time feedback.

Individual Difference	Description
Learning Objective	The aspect of technique that the player wants to improve through practice.
Playing Ability	An individual's current level of expertise on the violin and their strengths and weaknesses.
Approach to Playing	How they think about playing the violin and the way they like to practise. For example, whether they use mental visualisations.
Modality Preferences	Whether an individual prefers to use one sense over another.
Peripheral Awareness	How easy an individual finds it to be aware of objects in the periphery. How easily their attention is caught by exogenous cues.
Individual Motivators	The reasons that motivate an individual to practise the violin, such as a performance or for pleasure.

Table 8-4: Summary of individual differences that influence how real-time feedback is used.

Learning Objectives

Learning objectives are the elements of playing the violin that the individual player has identified as being problematic for them and are aiming to learn and improve. For example, Cathy in Chapter 7 had identified before the study that she needed to improve her bowing and was aiming to do this during the study. Feedback works most successfully when the individual's learning objective are the same as the focus of the feedback as this makes it relevant to the individual's own aims. For example, Cathy found the feedback very useful and successfully used it in her playing because the feedback focus was bowing which was the same as her own learning objective; whereas Beth ignored the feedback because she did not see bow length as a priority even though from an outsider's perspective it was something which she could have improved.

The relationship between feedback focus and an individual's learning goals affects taking action and catching attention in our model of using real-time feedback. Relevance to learning objectives will affect an individual's motivation to take action based on the feedback and their ability to be aware of the feedback. For example at the summer school in Chapter 7 Andy pointed out that not having an objective for using the feedback made it hard for him to notice it. Players using the feedback are having to dedicate some part of their attentional resources to peripherally monitoring the feedback in real-time. If they do not believe that the feedback is going to help them achieve their objectives then they will not dedicate the resources to it, making it harder for them to notice it.

Personal learning objectives will have an impact on motivation when using real-time feedback. As discussed in the literature review, there is a difference between intrinsic motivation and extrinsic motivation (deCharms, 1968) with intrinsic motivation being the most long-lasting. Whether real-time feedback should be classed as intrinsically or extrinsically motivating is down to the perspective of the player using it. For example, if a player has chosen to use it because she wants to use it to work on a learning objective which she has set herself then the way the feedback can be used to facilitate and measure her progress towards that goal will be intrinsically motivating because it is helping her to achieve something she is already intrinsically motivated to do. Conversely, if a player has been told to use the real-time feedback system and the system is giving feedback on something which the participant has no motivation to improve apart from to get better feedback then the feedback is acting as an extrinsic motivator. Although extrinsic motivation can be powerful, it can make an activity less intrinsically motivating (Greene & Lepper, 1974) and in the case of violin playing, which needs many years of dedication, feedback needs to support intrinsic motivation.

Learning objectives can be shaped by the learning context and recommendations from experts, for example the children from the high school study in chapter 7 took on improving posture as a learning objective because we spent several weeks discussing its importance and they were able to recognise how they needed to improve. Another example is chapter 5 when the participants used the learning objective set by the conductor of the orchestra out of respect for his authority. One way to make a design widely applicable to many people would be to give feedback to cater for many different learning objectives by giving multiple options for feedback focus. However, even more important than this is to support players in understanding their own playing and choosing appropriate learning objectives. Players who have teachers will already have access to support in this respect, but some players do not have regular teachers to help them recognise areas for improvement and form this into a learning goal which could be addressed with feedback. Therefore a feedback system needs to support players in analysing their own playing and choosing elements of technique to improve.

Playing Abilities

A player's strengths and weaknesses in the domain of playing will affect how much attention they can give the feedback because it controls task difficulty. How easy or difficult they find the task strongly affects how much they are able to notice the feedback. This was found in the studies in this thesis and in previous work with the MusicJacket (van der Linden, Johnson, Bird, Rogers, & Schoonderwaldt, 2011) where learners playing unfamiliar pieces had difficulty noticing the feedback. This shows how a task which is difficult as subjectively experienced by the player can prevent the feedback from being able to catch the attention of the player.

Approach to Playing

Individual approaches to playing also govern how real-time feedback is used. The way players think about their playing is important to the way real-time feedback works. Chapter 5 revealed differences in the way musicians think about the process of playing the violin. Some participants described visual imagery which they use to think about the action of playing. Another described using the written music to structure her analysis of her own playing. Others were drawn to tactile feedback because they felt that it fitted with process of playing, showing that they must think about playing in terms of tactile and kinaesthetic sensation. Feedback that fitted with the particular way an individual thought about playing the violin was described as being easier to understand than other types of feedback by that individual. This shows that links between the modality and metaphors used in the design of the feedback and an individual's own ideas about playing can help a player to interpret the feedback.

Modality Preferences

The individual approaches to playing that the participants described and some of the reasons for preferring feedback in a particular modality suggest that this may stem from an innate preference for using and thinking with a particular modality. This preference appears to be stronger in adults than in children. There are several possible reasons why individuals may have a preference for using a particular modality: first it may appeal to the way they like to think and learn, second they may be better at perceiving stimuli in that modality, third, they may find that modality more comfortable to use.

Some participants in the modality comparison study in chapter 6 chose one modality over another because they found it fitted with the modality that they already used for thinking about the violin playing the violin. This may be linked to the concept of a learning style. There are a number of theories which aim to categorise people in terms of learning style, in particular Fleming's VARK model (Fleming, 2006) which categorises people as Visual, Auditory, Reading/writing and Kinaesthetic learners. A learning style is an individual's preference towards learning and thinking in a particular way and most advocates of learning styles argue that if learning materials can be tailored to the learning style of individuals' then they will learn more efficiently. In the literature there is general consensus (Pashler, McDaniel, Rohrer, & Bjork, 2008) that it is common for individuals to express a preference towards learning using a particular modality or set of modalities. However, there is controversy over whether tailoring teaching to this improves learning (Pashler, McDaniel, Rohrer, & Bjork, 2008). Learning styles may be more relevant to real-time feedback for music training than for traditional classroom learning as real-time feedback needs to be perceived and understood on a much shorter time scale (within seconds). Therefore, preferences toward using a particular modality or learning style may have a more substantial effect on the learning outcomes because if a player is able to understand the feedback stimulus half a second faster in their preferred modality then this will have large impact on the effectiveness of the feedback. Further research is needed to investigate whether there is a connection between learning style and the most effective modality for giving real-time feedback.

There are also individual differences that stem from individual's perceptual abilities. There is variation in people's abilities to perceive vibrations (e.g. (Halonen, 1986)). Age and gender can both influence the threshold at which people are sensitive to vibrations, but more than this there is considerable variation which is not accounted for by either of these two factors. Similarly people differ in their visual perception and their ability to use their peripheral vision in different situations (e.g. (Williams & Andersen, 1997)).

Another reason for preferences towards using a particular modality is related to comfort. Although most participants did not find the vibrations uncomfortable, one participant in Chapter 5 found the vibrations extremely unpleasant and others said they found the sensation annoying. This is supported by other studies which show that there is variation in people's perception of what is comfortable when it comes to vibrations (Nagel, Carl, Kringe, Märtin, & König, 2005), (Pielot & Boll, 2010). A large scale study (Neely & Burström, 2006) found that women have a lower threshold at which they find vibrations uncomfortable than men. Issues of comfort are not confined to vibrotactile feedback; there was also an individual in the study with MuSense who found that having flashing in her peripheral vision was uncomfortable.

In sum, people commonly expressed a preference towards feedback in a particular modality. Comfort, perception and learning styles are all potential reasons for this. However, most people were able to use non-preferred modalities as well and it is not yet clear whether the preferred modality gives better performance or simply a better experience using the feedback. Further research would be needed to confirm this.

Peripheral Awareness

People differ in their level of peripheral awareness when concentrating on demanding tasks. This affects how real-time feedback catches the attention of the player. As already discussed, using flashing lights or other sudden changes makes feedback more noticeable, but how this conspicuous signal is experienced is dependent on the individual. In the modality comparison study in chapter 6, some participants found the flashing lights too noticeable, intruding into their main focus of reading music; in contrast some found the visual feedback was not obvious enough. This shows that there are individual differences in people's peripheral vision which impact on how effective different visual feedback designs will be for that person.

Other studies have also found variation in people's peripheral vision while carrying out demanding tasks. In sports science, Williams and Andersen (1997) have studied athlete's response to stressful situations and found that vision narrows in demanding situations. They showed that the reduction in peripheral vision is related to psychosocial factors such as negative life experiences (Williams & Andersen, 1997). Playing music is also demanding, and being observed in a study it may be particularly stressful. Therefore it is likely that participants are also exhibiting similar stress responses which narrow their peripheral visual field. The variation in individual awareness of peripheral visual feedback may be due to similar causes to those found in the athletes which Williams and Andersen (1997) studied.

Individual Motivators

The study with the high school students in Chapter 7 showed that an individual's sources of motivation and the type of goals they adopt affects how the view real-time feedback. This makes sense because goal orientation (e.g. whether players practice for mastery, or aims to perform well in front of others, or avoid playing badly in front of others) affect the way people practice without feedback (Smith, 2005). The discussion in Chapter 7 identified five difference sources of motivation that may influence how real-time feedback is used: aesthetic pleasure in playing, intellectual curiosity, a sense of achievement, performance goals and competition.

These motivators interact with characteristics of the feedback and the external factors to determine whether a learner is motivated by the feedback.

Aesthetic pleasure in playing is an important intrinsic source of motivation for playing the violin. Feedback that adds to the aesthetic experience and expressivity of making music such as attractive visual displays or even additional music as feedback may enhance this motivation. This can also be seen in the way Andy liked using the visual feedback to augment and add additional expressivity to his performance at the summer school in chapter 6. In contrast, feedback which is perceived as *forceful* may damage this sense of intrinsic motivation as it is removing the locus of control from the learner.

Intellectual curiosity in finding out new things about their playing and music in general was a source of motivation for the children in Chapter 7. Relatively *complex* feedback that allows exploration of movement and enables a new perspective on playing may add to this source of motivation, whereas feedback that is *forceful* or too simplistic would not.

Some people are motivated by a sense of achievement which may come from mastery goals which focus on mastering the instrument. For these children, feedback that explicitly shows them how they are progressing and allow them to focus on their achievements will be motivating. Conversely, strongly *negative feedback* that dwells too much on mistakes without giving an opportunity to see progress would be demotivating.

Others are motivated by performance goals which are concerned with how others view their playing. With this goal orientation practice may be seen as a functional activity purely to prepare the player for performance. In this case *forceful* feedback that gives a sense of being made to improve may be preferred. Players may also be motivated by competition. This relies on a *social context* where people are playing together, are equally matched and have the type of relationship where competition can flourish. In this situation feedback, such as the Twinkly Lights in Chapter 7, that is *publicly visible* will facilitate competition. Further researcher is needed into how real-time feedback can be designed for people's individual sources of motivation and whether this enhances the effectiveness of feedback or the experience of using it.

Individual Differences – Summary

Individual differences can have a large influence on the way a person experiences feedback. Although many approaches to research try to show a single design that works for all people, from the perspective of the individual using the technology it is better to cater for individuality, particularly because different designs of feedback can result in such differing experiences as described in the comparison study in Chapter 5.

8.4 Discussion

The framework in this chapter described how visual and vibrotactile real-time feedback can be designed as a technology intervention for improving motor skills learning, based on the findings from the studies in this thesis. The in-the-wild nature of the research, combined with the qualitative and participatory methodology has also led to a focus on the individual user rather than users as group, which has uncovered many individual differences. Because of this methodology, a rich and detailed framework has been produced that covers many factors that influence user experience and use of real-time feedback. However, by covering such a wide area, some of the interactions between factors are not yet fully described raising many questions for future research. For example the details of personal motivators that influence the use of feedback was only discovered in the final study and further research could uncover how these differ between adults and children and whether manipulating characteristics of the feedback could support these different sources of motivation.

The studies in this thesis relate to the physical aspects of music learning. However, there is potential for the framework to be generalizable to other types of motor-learning. Music learning is very similar to other motor-skills such as learning sports or good posture or physical therapy exercises. These also require motivation, guidance about movement and sustained practice. External factors such as task, social context and learning context will also play an important role in the suitability of different feedback designs and these will interact with the characteristics of the individual. For example, when analysing a design to encourage good posture at work, social context will play an important role, it may act to make people feel self-conscious about any type of visual feedback; alternatively in the right kind of office environment visual feedback may stimulate competition. The focus of the feedback on posture and how this relates to individual learning goals regarding posture will also be important. It may be possible to encourage a positive attitude towards posture goals by introducing the feedback as part of a learning programme about posture to give it a favourable learning context. Individual differences will influence which feedback characteristics are most effective for different workers, e.g. peripheral awareness will interact with the conspicuousness of the feedback to decide whether it is useful for improving posture or whether they do not notice it or find it distracting from their work.

In different applications different factors will dominate and some new factors may be discovered to extend the model. For example, when designing real-time feedback to encourage exercise in those with chronic pain (e.g. Singh et al., 2014), individual's motivators will be particularly important and will be different to the ones that motivate people to play music. Positive feedback may be more advisable as negative feedback for because people with chronic pain may easily be discouraged, or their anxiety increased. Whereas, in music playing players can use negative feedback to measure their progress because they are supported by strong motivation and self-confidence. Self-confidence may be a new factor under individual differences that would need to be introduced if this framework was applied to real-time feedback for chronic pain.

Currently this framework only describes real-time vibrotactile and visual feedback. This is a constraint imposed by the feedback used in the studies in this thesis. Another key area where this research could be extended would be to investigate other modalities. However, this should be for a different application because music learning is not well suited to auditory feedback.

Chapter 9 – Guidelines for Designing Peripheral Real-time Visual and Vibrotactile Feedback

9.1 Introduction

This chapter outlines some practical implications for the design of real-time feedback in the form of guidelines that arises from the findings of the studies conducted in this thesis and the theoretical framework developed from them. Although the focus of this thesis has been on providing real-time feedback to aid learning the violin, it is proposed that the guidelines are more generally applicable to real-time feedback for learning any motor skill – particularly if the skill is demanding perceptually and cognitively. The guidelines are structured in terms of the feedback characteristics that they relate to based on the real-time feedback framework. Guidelines concerning each characteristic are grouped together and then discussed with respect to the evidence from studies supporting them and their applicability to other domains.

Throughout this thesis it has been recognized that core design issues come from understanding the demands of the context which the feedback is being used in. Understanding the effect of context on how prototypes are used and how different contexts place different demands and stresses on users through in-the-wild studies has been essential for moving the design of the prototypes forward. Here we discuss guidelines for designing real-time feedback for the different learning contexts that we investigated: individual practice and ensemble practice. Another aspect that became an increasingly central concern in the programme of research was calibration. It was found that designing for calibration evolved through the design iterations, until it became an essential part of learners understanding how the feedback worked. Here we discuss how to design calibration so that users can understand it and why this is so important.

9.1 Guidelines

Table 9-1 show the guidelines for designing visual or vibrotactile real-time feedback. These have been structured into categories corresponding to the feedback characteristics discussed in the last chapter. These characteristics are: Feedback Focus, Modality, Position, Conspicuousness, Complexity, Public visibility, Metaphors, Positive versus Negative Feedback and Forcefulness. In addition there is the Context category. For each of these categories the guidelines are individually discussed and explained in more detail with particular consideration for their applicability in other domains. At the end of this chapter guidelines about calibration are presented which are generalizable beyond real-time feedback applications to other sensing applications.

Factor		Key Guideline	Specific Guidelines
Feedback Focus		Choose the feedback focus with the learning context in mind	Fit feedback into a programme of learning.
			Match the feedback focus to people's learning objectives.
Modality		Visual nor vibrotactile can both be effective neither is superior	Take into account individual differences in how people respond to different modalities
			Choose a multimodal system rather that communicating too much information in one modality
Position	Vibro	Position vibrators with respect to the anatomy underneath the skin	Position vibrotactile feedback near to reference points (e.g. joints)
	Visual		Position peripheral visual feedback near to the main visual focus
Conspicuousness		Fit the level of conspicuousness to the context and the individual	Use sudden changes in feedback to catch attention
			Take into account the demands and stresses of the context.
			Take into account individual differences in how aware people are of their peripheral vision.
Complexity		Choose a single feedback focus	Keep peripheral visual feedback very simple
Public Visibility		Take into account social context when design publicly visible feedback	Visual feedback is publicly visible whereas vibrotactile feedback is not
			Design public/shared feedback to motivate groups to support one another
			Competition can motivate people, but this depends on the context
			Still position visual close to each individual's main visual focus but also where others can see it.
Metaphor		Use embodied metaphors to describe feedback	Before designing visual feedback research the way learners and teachers think about the movements they are learning and capitalise on these
			Aim to reflect body movement in real-time through movement in visual feedback.
			Consider how multiple metaphors may be used in parallel to offer a variety of ways of thinking about feedback.
			Do not rely on a metaphor to be a full explanation of the feedback.
Negative vs Positive		Both negative & positive feedback is effective	Aim to minimise how much the vibrotactile feedback is triggered
			Certain visual signals can have positive or negative connotations that can be exploited
Forcefulness		Feedback can be designed to be forceful or optional	Take into account individual differences in motivation when designing forceful feedback
Context		Take into account the demands and stresses of the setting	Carryout user testing real-world settings with the same demands and stresses as the setting being designed for.
			In lessons, design feedback to facilitate communication between teacher and pupil
			In individual practice, design feedback to support self-analysis
			In ensemble practice, design feedback to encourage mutual support

Table 9-1: Table of guidelines for designing real-time feedback

Feedback Focus

Guidelines	Chapters	Factors and issues
Match the feedback focus to people's learning objectives.	5, 6, 7	Individual differences, motivation, taking action
Fit feedback into a programme of learning.	7	Motivation, ease of interpretation, taking action

As the research has shown, feedback has an element of individualisation that makes the feedback suited to the learner's level of playing and their particular playing problems. At a fundamental level, the focus of feedback needs to match a student's learning objectives. This can be designed for in two ways. Firstly, a number of different options for feedback focus could be given so that players can choose the focus that fits their learning objectives best. Secondly, additional support can be given so that feedback is given as part of a programme of learning as in the study with the high school students in Chapter 7. By making the feedback part of a programme of learning the students can be taught to value posture and analyse how they can improve their own posture before having the feedback.

The relationship of the feedback to learner goals is relevant to any feedback application. For example, if designing feedback to help somebody's golf swing a system would need to be able to give feedback about the specific element of their swing that is problematic for them. Alternatively, if the golfer was a novice who did not have the knowledge to identify problems with their swing, the feedback could be part of a programme of learning. As the player gradually learns about swinging the club and what is desirable in a swing and why, they could then use the feedback to help them to diagnose and work on problems.

Modality

Guidelines	Chapters	Factors and issues
Of the visual and vibrotactile modalities both modalities can be effective and neither is superior to the other	5	Individual differences, Context and Task
Take into account individual differences in how people respond to different modalities	5	Individual differences, catching attention, taking in information, interpretation
A multimodal system is preferable to too much information communicated in one modality	(Johnson, van der Linden, & Rogers, 2010), 7	Catching attention, taking in information

Individual selection The studies showed that on average both types of feedback were equally effective. However, on an individual level there were marked differences in how people responded to each modality. This means that when designing feedback either modality can be used effectively but the best design might offer an option to use either modality based on which is most effective for the individual. Some individuals may be able to identify accurately which modality is best for them but others may not. Therefore when designing for different modality options a method that enables the user to find their best modality is desirable.

Multimodality It is possible to have two kinds of real-time feedback that correspond to two elements of a skill. For example, if designing a system for physical therapy that is intended to encourage both good posture and a large range of movement in the arms, then vibrations on the body could be used to encourage and upright position, while a visual display is used to reward reaching the arms outwards.

Although both can be effective, vibrotactile and visual feedback are different, and this should be taken into account in how they are designed. Therefore some guidelines in the following sections are specific to one modality or another. There will be also situations when one modality is preferable to another, e.g. in an activity that involves intensive body movement it may not be possible to feel vibrotactile feedback (Post, Zumpa, & Chapman, 1994). Similarly, if the activity involved intensive head movements then this may prevent a person focusing on visual feedback.

Position

Guideline	Chapters	Factors and Issues
Position vibrators with respect to the anatomy underneath the skin	3.7	Comfort, Clarity of perception, Taking in information
Position vibrotactile feedback near to reference points (e.g. joints)	3.7	Clarity of perception, Taking in information
Position visual feedback near to the main visual focus (e.g. the music stand for violin playing)	5, 6, 7	Catching Attention

Vibrotactile Feedback

The specific anatomy directly underneath the vibration motor plays an important role in the comfort of the vibrations and how easy they are to differentiate between when under cognitive load. Vibrators positioned on bone produce a different sensation to those positioned on muscle. Research designing the MusicJacket suggested that vibrators on bone were more comfortable (Appendix A). On the other hand, bone conducts vibration which may make the location of the vibrator difficult to perceive. Torso vibrators positioned near the diaphragm can give an uncomfortable sensation, whereas on the ribs they give a different sensation which is more comfortable.

Peripheral real-time feedback also needs to give clear signals as users are only able to pay a small amount of attention to the feedback while doing other tasks, therefore it is best to position vibrators close to the reference points on the body such as joints (Cholewiak & Collins, 2003) so that they can be easily identified. This is different to many designs in the literature (e.g. (McDaniel, Goldberg, Bala, Fakhri, & Panchanathan, 2012), (Spelmezan, Jacobs, Hilgers, & Borchers, 2009)) which use a denser arrangement of vibrators on the body, but these are designed for user's focused attention, whereas our lighter displays are better suited for peripheral feedback.

The research in this thesis only studies vibrators positioned on the arms and sides of the torso, which may limit the applicability of the findings to other parts of the body. Nonetheless, there is a strong argument that if the anatomy beneath the skin is important in these parts of the body it will also play a role elsewhere. For example, if designing an application for skiing that gave vibrotactile feedback on the legs, first it would be important to think about the anatomy of the legs in terms of what types of tissue are near the surface of the skin and test the different in sensations produced by positioning vibrators on different types of tissue on the leg.

Visual Feedback

It is best to position peripheral feedback near to the main visual focus for the task (Lamble, Laakso, & Summala, 1999). For music learning positioning the feedback on the music stand near the music is the best position for catching attention. This is because a player's main visual focus is on the music and by positioning the feedback close to this and at the same focal distance away, participants can switch their attention to the display more easily. For example if designing visual display to encourage good posture while working on the computer, it would be best to mount it on the monitor or have it as part of the display on the monitor. In other domains such as sport, this guideline may be more difficult to implement as in sport a person's visual focus may shift constantly. Simple feedback on HUD displays may work here for applications such as skiing. For practicing specific actions in sport such as shooting a basket with good technique in basketball, the visual display could be mounted onto sports equipment such the backboard of the basketball net.

Conspicuousness

Guideline	Chapters	Factors and Issues
Take into account the demands and stresses of the context and task when thinking about conspicuousness.	5, 6, 7	Context, perceptual and cognitive load
Use sudden changes to catch attention	5, 6, 7	Catching attention
Take into account individual differences in how aware people are of their peripheral vision.	5, 6	Individual differences, comfort, catching attention

Vibrotactile Feedback

The conspicuousness of vibrotactile feedback in our designs was mainly dependent on the positioning of the vibrators therefore we have not added any additional guidelines for vibrotactile feedback concerning conspicuousness. There may be ways to adjust the conspicuousness by adjusting the intensity of vibration and the way in which they are activated (e.g. patterns of vibrations). However, these were not investigated here as we found a simple on-off activation at a medium intensity to be sufficient. It is also important to note that one of the largest influences on people's awareness of vibrotactile and visual feedback is the setting and difficulty of the task. Thus it is important to take in to account the demands and stresses of the context of use when thinking about whether a design will be noticed by users.

Visual Feedback

Sudden changes in a visual display are better at catching attention than continuous ones (Mulckhuyse & Theeuwes, 2010). Chapter 5 showed that flashing is a good way for feedback to catch the attention of the player. However, flashing can be obtrusive and difficult to ignore. To address these differences in the way visual stimuli catch people's attention, players could be given options about how conspicuous the feedback should be to suit their needs. Players who find it difficult to notice the visual signals could choose feedback which uses flashing to catch their attention; players who find the flashing draws their eye away from the music too much could choose to have this part of the display switched off.

When designing visual feedback for other settings, for example, improving posture at a computer, a flashing display will be better at attracting the attention of a worker who is prone to becoming engrossed in their work. However, a different person may be too easily distracted by such a display and it would be better to give them the option to have a less conspicuous display. There will also be different requirements for the conspicuousness of the display depending on the type of work being carried out and how engaging or demanding it is. In a shared office environment a visual display may have to be made less conspicuous due to the social context.

Complexity

Guideline	Chapter	Issues and Factors
Choose a single feedback focus for vibrotactile feedback	3, 5, 6, 7	Taking in information, reducing cognitive load
Keep peripheral visual feedback very simple.	5, 6, 7	Catching attention, taking in information
Choose a single feedback focus for visual feedback	5, 6, 7	Taking in information

In applications such as playing the violin where the primary visual focus is not on the feedback it is important that visual feedback is kept simple. The display used in the summer school study in chapter 6, although working well when participants gave it their full attention, was too complex to take in while playing. Even the moving line display (complex lights) on the MuSense prototype was considered too complex by some people to take in while reading music.

This suggests that a display showing one dimension of movement is the most that participants can attend to while playing. When designing such a display it is important to make sure that details such as colour or number are not vital to interpretation compared to the overall form or gestalt which is seen at a glance.

Vibrotactile feedback also works best when limited to a single feedback focus at a time. This gives clarity to the feedback signals by reducing complexity and allows sparse positioning of vibrators. In early studies with the MusicJacket, where vibrotactile feedback was given about several elements of playing at once, players became overloaded (summarised in Chapter 3). Research in the literature that has used more complex vibrotactile displays with patterns of vibrations suggests that people are able to take in this complexity in a lab when it is their main focus e.g. (McDaniel, Goldberg, Bala, Fakhri, & Panchanathan, 2012), but our research would suggest that when vibrotactile feedback is being used in real-time in real-world situations users have difficulty differentiating between signals.

The complexity of feedback needs to be matched to the demands of the setting and task. For example, if designing feedback to guide shooting a basketball at a net while trying to be aware of attacking players a single feedback focus for the vibrations would be best, whereas if it is for individual practice with no other distractions slightly more complex feedback maybe possible.

Public Visibility

Guidelines	Chapters	Factors and Issues
Take into account that visual feedback is publically visible whereas vibrotactile feedback is not	6, 7	Modality differences
Take into account the social context when deciding whether to design a public display	6, 7	Context, individual differences
Design public/shared feedback to motivate groups to support one another	6, 7	Motivation, Taking action
Competition can motivate people, but this depends on the context	6, 7	Catching Attention
Still position visual feedback close to each individual's main visual focus but also where others can see it.	6, 7	Motivation, Taking Action, Context

A key difference between visual and vibrotactile feedback is that visual feedback is publicly visible whereas vibrotactile feedback is not. Social context is particularly important when considering whether to design publically visible feedback. All the groups in this research were supportive of one another and we found that publicly visible feedback could be motivating in this environment. For example in Chapter 7 the high school children suggested using the visual display as part of a game because they found competing with one another motivating. The public nature of the display also meant that if one child was not paying attention to the feedback the other would remind them.

It is important to position shared visual feedback near to the main visual focus. We found that having a shared projected display was too difficult to use because it did not catch people's attention; whereas, the children's individual Twinkly Lights on the music stand worked well because they were near to the music and visible to others. Visual feedback on the stand also shows clearly who the feedback belongs to and it naturally prioritises the player to notice their own feedback first (because it is closer) before looking at others. The high school study in Chapter 7 found that it is not necessary design competition into the feedback itself, so long as players are able to make a fair comparison between themselves and their peers. Competition could be enhanced by giving additional ways for participants to make this comparison, e.g. giving extra statistics about playing at the end of a piece. However, this is specific to the setting, in a less supportive environment or where individuals have lower self-confidence competition would be inappropriate.

Publically visible feedback can be useful for many applications but may be inappropriate for some. For example, public feedback could be good for team sports training because the team needs to learn to work together and the context is competitive so being publically judged is appropriate. In contrast, physical rehabilitation would not be a good setting for publically visible feedback because each person is making their recovery journey individually and they may be vulnerable and have low self-confidence. They may also face setbacks that are beyond their control and it would not be appropriate for public feedback to encourage others to make judgements about their performance. However, there may be some rehabilitation groups that are very supportive of one another and some kind of public feedback that helps people show their support for one another might be appropriate if designed considerately.

Metaphor

Guidelines	Chapters	Factors and Issues
Before designing visual feedback research the way learners and teachers think about the movements they are learning and capitalise on these	5	Ease of interpretation
Aim to reflect body movement in real-time through movement in visual feedback.	5, 6, 7	Ease of interpretation
Consider how multiple metaphors may be used in parallel to offer a variety of ways of thinking about feedback.	5,6,7	Individual Differences, Flexibility
Do not rely on a metaphor to be a full explanation of the feedback.	5, 6, 7	Ease of interpretation

Vibrotactile Feedback

As discussed in the previous chapter having a good metaphor helps participants understand how to interpret the feedback. Vibrotactile feedback is positioned on the body and lends itself well to embodied metaphors draw on the everyday ways we interact with our environment. The metaphor used in the MusicJacket (van der Linden, Johnson, Bird, Rogers, & Schoonderwaldt, 2011) and Buzzy Jacket (Chapter 7) was the *push* metaphor where vibrators were positioned on the arm to push the arm into the correct position when a mistake is made. This metaphor was helpful in explaining what the vibrations meant, but it is an incomplete explanation because there are multiple possible responses that fit with the push metaphor but only one is correct from a violin playing perspective. Therefore a designer should not rely entirely on a metaphor to explain the feedback. It is important to consider how the metaphor could be misconstrued and what support can be built into the design or may be present in the learning context (e.g. teachers) to ensure people respond to the feedback correctly.

Another embodied metaphor is the idea of moving the body into *different zones* e.g. the elbow on the Buzzy Jacket vibrates because lowering the violin causes the elbow to move down into an undesired zone. This metaphor can work in parallel with the push metaphor because the positioning and activation of the vibrators are consistent between the two metaphors. This shows that for some set-ups more than one metaphor can be used to explain what the feedback means. These are useful because they can provide more scope for how to design and use real-time feedback.

Zones were also used in the design of the feedback for long bowing; in this case the idea of getting vibration for reaching a *good* place. In some ways metaphors about moving the body into *different zones* are more versatile than the *push* metaphor because they can be used for both positive and negative feedback.

How to apply a metaphor depends on the action being practised. For example, vibrotactile feedback could be used to encourage follow through in a forehand tennis swing by giving a buzz when the participant's bat enters a zone near to the left shoulder, showing that they have done the correct follow through.

Visual Feedback

Visual feedback suits different metaphors compared with vibrotactile feedback. Vibrotactile feedback can point to parts of the body using a first person reference frame but is low resolution; whereas visual feedback cannot reference parts of the body in such a direct way but is higher resolution and can reflect body movements through movements in the feedback display. Some people already use visual metaphors to think about their playing, for example the university orchestra players using MuSense explicitly described some of the metaphors they used to think about their bow, such as visualising it as a line. If visual feedback is designed to capitalise on the metaphors people already have for imagining their playing then this will help them to interpret the feedback. Therefore, the design of real-time feedback can be improved by building upon existing metaphors used by potential users. Moreover if there are several common metaphors used by people it may be worth trying to design the feedback so that it is consistent with all these ways of thinking, or allow people to select their preferred metaphor.

Metaphors are also used in teaching to encourage particular movements and these could be enhanced through feedback. For example, in the preparation sessions with the high school children the metaphor of a thread pulling from the top of the head to the ceiling was used to encourage the pupils to sit-up straight. This could be used to design visual feedback to encourage good upper body posture.

Although there is a lot of potential for metaphors in visual feedback, caution should be taken not to let metaphors become too complex so that it can still be effective peripherally in real-time. If it becomes too complicated it can disrupt the on-going skill, e.g. the flower metaphor used in the ensemble was too complicated for the players to interpret while playing.

The above guidelines are widely applicable to visual feedback in any domain involving movement. For example, when designing feedback to encourage skiers to keep their legs bent, it would first be best to interview skiers to find out how they visualise their knees while skiing and interview their teacher's to find out how they describe the ideal movement and any metaphors they use. This could then be used as a basis for the design of the visual feedback. If the knees are supposed to be springy then springs or other representation of springiness could be used as a visual metaphor.

Negative and Positive Feedback

Guideline	Chapters	Factors and Issues
Negative and positive feedback setups are can both work effectively irrespective of modality.	5, 6, 7	Motivation,
The main issue is to minimise how much the vibrotactile feedback is triggered.	5	Avoid perceptual overload
Certain visual signals have positive or negative connotations. These can be exploited in the design of real-time feedback.	5, 7	Ease of interpretation

Vibrotactile feedback was used in this thesis both positively to indicate a goal achieved and negatively to indicate a mistake. The interviews in Chapters 5 and 7 suggest that the amount of time the feedback on should be minimised. So, for long bowing it makes sense to give a signal at the end of the bow to indicate reaching far enough, rather than having it on in the middle and having it switch off to show they had reached a good length. When participants were asked to come up with other designs they switched between positive and negative feedback depending on the design so that the vibrotactile feedback would be on the least.

Visual feedback has pre-existing positive and negative connotations (e.g. the green lights on MuSense were agreed by participants to be a positive signal). However, design does not have to be constrained conventions for what form positive or negative feedback should take (e.g. colourful lights were used successfully in the high school study to indicate mistakes in the children's posture even though Christmas lights have positive connotations).

To summarise, both vibrotactile and visual feedback are suitable for both positive and negative feedback and neither setup is inherently demotivating. The choice of whether vibrotactile feedback is positive or negative should be based on which one will trigger the least feedback and makes most sense in terms of the other metaphors being used. For example, when designing real-time vibrotactile feedback for physical therapy to encourage patients to increase their range of movement through stretching, it would be best that the feedback is triggered as a reward when they reach out far enough to stretch their range of movement rather than being used as a negative indicator that they have not gone far enough as this minimises the time the vibrations are switched on. Visual feedback is flexible because it can be designed to have positive or negative connotations, it will depend on the feedback focus whether it is better to draw attention to mistakes or encourage users towards a goal. In the stretching example it would be best to use the visual feedback as a reward because extending range of movement is a goal to be reached for rather than their current movement being a bad habit to avoid.

Forcefulness

Guideline	Chapters	Factors and Issues
Feedback can be designed to be forceful or optional	7	Taking Action
Take into account individual differences in motivation when designing forceful feedback	7	Individual differences, Motivation

Some designs of feedback are perceived as insisting upon an action in response, whereas other feedback is seen as information only. The vibrotactile feedback on the Buzzy Jacket was perceived as being forceful whereas the Twinkly Lights was not. Nonetheless, it may also be possible to design visual feedback that is forceful.

When choosing whether to design forceful feedback or not it is important to consider individual differences in motivation. Players that are motivated to improve for an external goal such as a performance may appreciate forceful feedback because it will help them improve quickly. Players that are motivated by the intrinsic pleasure of playing may dislike forceful feedback because it detracts from the experience of playing. Similarly in sports, those training for a competition may choose to use forceful feedback to improve their technique; but those who practise sport for pleasure may prefer something that enables more choice.

Context

Guidelines	Chapters	Factors and Issues
Take into account the demands and stresses of the setting being designed for	5, 6, 7	Context, cognitive load, motivation
If informally testing a design (self-testing or similar) still try to test it in the target setting on similar tasks	6	Context, cognitive load
Carryout user testing real-world settings with the same demands and stresses as the setting being designed for	5, 6, 7	Context, cognitive load, motivation
In individual practice, design feedback to support self-analysis	5, 6, 7	Ensuring correct calibration and interpretation, no learning support from context
In ensemble practice, make feedback as simple as possible	6	Catching attention, ease of interpretation, reducing cognitive load, demanding and high stress setting
In ensemble practice, design feedback to encourage mutual support	6, 7	Motivation, catching attention

Context is central to how feedback is used; playing in an ensemble is a completely different experience compared to playing at individually. Thus feedback design needs to be matched to the demands of the context. It is important to stress that whatever setting is being designed for, the prototype should be tested in that setting. It is common for designers to test prototypes on themselves while they are in the process of building them. Even at this stage it is important that a designer should try to simulate the environment the prototypes will be used in. For example, if building it for individual practice the designer should test it while playing a difficult piece of music and test it for a full length practice; or if it is for ensemble playing, the designer could try to get a colleague to play with them so that they are feeling the type of pressures users will feel. When a full prototype has been built, user studies need to then be conducted in realistic settings because this is where the most revealing findings for design can be discovered.

Individual Practice – Support Self-Analysis

Individual practice is a simpler setting than ensemble practice setting to design as there are fewer demands from the social context and the environment. However, this can also be a drawback because there are no experts or peers to help or prompt the players to use the feedback correctly. Therefore, an additional design criterion that needs to be addressed for individual practice is how to promote learners to analyse their own playing and set appropriate learning objectives.

Feedback in real-time cannot make a person understand every detail of their movement – the bandwidth of what people can take in is too low. Nor can it tell them exactly how they ought to be playing – each player is too individual and there is no consensus on how the violin should be played, this is why we introduced calibration processes. In the studies there were occasions when participants were not aware of the problems with their playing, and using the real-time feedback did not always help the player to notice or accept these problems. In individual practice there is nothing to stop people focusing on the wrong learning objectives or misinterpreting the feedback to fit with their pre-existing ideas about their own playing. A teacher can support players in overcoming these problems and help them to set goals. However, as a player matures they must also learn to do this for themselves through self-analysis. Feedback designed for individual practice needs to support the player in analysing their own movement, with or without the help of the teacher, and from this to set appropriate goals. These can then be implemented through the process of calibration and real-time feedback can be used as a tool to achieve these goals through practice.

How to go about supporting self-analysis is a difficult design challenge. It may be necessary during the goal setting stage to give more comprehensive feedback after the learner has finished playing so that they can recognise areas they would like to improve. Existing methods that some teachers already use such as using mirrors and lesson and practice diaries could also be combined with this to aid the process of reflection.

Ensemble Practice – Simple Feedback and Mutual Support

When designing for ensemble practice the most important design criteria are simplicity and capitalising on group dynamics. Ensemble practice is a particularly cognitively demanding setting this is why a simple feedback design is so important in this setting. There are different kinds of ensemble practice which have different cognitive demands. The most difficult was the kind of ensemble playing studied at the summer school where the player has to hold an individual part against other members of the group and the music is unfamiliar. The main aim of those ensemble practices was to learn the notes and to learn to play them in time with one another. This made using the feedback challenging due to high cognitive load and a mismatch between feedback focus and learning focus. In contrast, the children at the high school were all playing the same part and had time to become familiar with the music making their task easier. There was also joint aim for the whole group to improve posture. This meant that real-time feedback was well suited to this situation. Nonetheless even in this situation ensemble playing is more demanding than individual playing and it is important the feedback should be simple and easy to take in.

Ensemble practice presents opportunities to draw on the dynamics of the group to help motivate players to use the feedback. Publicly visible feedback, such as the Twinkly Lights on the music stand in the high school study, is more effective for stimulating group discussion than invisible feedback like vibrotactile feedback. Such feedback can be used by a group to compete with one another and prompt group members to be explicit about the feedback if they had not noticed it. Both these mechanisms are beneficial for encouraging the use of real-time feedback. Therefore, real-time feedback for group settings, such as ensemble playing, should aim to encourage mutual support and friendly competition.

Applicability

The specific guidelines outlined about the different settings studied in this thesis (individual practice and ensemble practice) will be most relevant to those designing for musicians. However, they may also crossover into other domains for example they may be useful for thinking about the differences between designing for individual sports practice and team practice. The key message here is to learn about the demands and stresses that people experience in the learning setting and to account for these in the design.

9.2 Calibration

Guidelines	Chapters	Factors and Issues
Use calibration as a way for learners to set goals	7 and MusicJacket (Johnson et al., 2013)	Individual difference, motivation, user control
Design calibration to be appropriated	7	User control, motivation
Expose intentions and provide visibility about how the system functions and what it is aiming to do.	7	User control, appropriation
Use familiar mappings such as drawing around the body to take measurements.	7	Ease of interpretation, user control
Design for learning through tweaking.	7	Ease of interpretation, getting a good calibration, user control
Contextualise the calibration process in relevant domain knowledge relating to the feedback focus.	7	Getting a good calibration, user control

Feedback must be appropriate to the level, body shape and playing style of the learner. This can be achieved through calibration. Initially, we viewed calibration purely as a requirement of the technology to have a baseline from which to make judgements about good and bad playing. However, the research in this thesis has shown that how people understand calibration is important to how the feedback is used. Calibrating an ideal baseline is really a form of setting a goal. Participants need to engage with this so that they take ownership of the goals the feedback is supporting.

The findings concerning calibration from the user studies show that calibration can be more than just fulfilling a functional requirement of a particular technology. The Buzzy Jacket showed that it can be used by students to set their own goals and give them the flexibility to play in a way they feel comfortable with, rather than fitting a rigid model. The Footboard and Twinkly Lights demonstrated that the process of calibrating can also become part of learning something new and useful. Earlier research with the MusicJacket has also shown that teachers can use it to set learning goals with students (Johnson, Bianchi-Berthouze, Rogers, & van der Linden, 2013).

In the kind of calibrations discussed in this thesis – namely calibrations which individualize the system to the particular needs of the user – helping users to engage with and understand the calibration process was found to be beneficial. Calibration can potentially give users control over the key functions of the system, e.g. calibrating the Buzzy Jacket let them control where the vibrations were triggered. However, designers often view calibration only as a technical process and place little emphasis on helping users understand it. By de-mystifying calibration we can enable users to appropriate it and take ownership of what the system is doing. In the field of education this is particularly important because taking ownership and responsibility for learning is key to learner development. In the box below and on the next page are guidelines concerning how to design calibration that is easy for students and teachers to understand and helps them to learn about how the system works.

9.2.1 How to Design Calibration that is Easy to Understand and Appropriate

Here are guidelines for designing calibration that is easy to understand and appropriate, based on (Johnson, Bianchi-Berthouze, Rogers, & van der Linden, 2013). Depending on context appropriation will be by the learners themselves or by their teachers.

Expose intentions and provide visibility. Similar to Dix (2007), we argue that designers should think about the purpose of calibration for the applications they are building. In order to calibrate well and get the benefits from engaging with calibration the users need to understand the purpose of calibrations (the designer's intentions) and how calibration affects how the feedback works afterwards. For example, the intention of calibration became clearer to the violin teacher in the MusicJacket study (Johnson, Bianchi-Berthouze, Rogers, & van der Linden, 2013) through discussion with the researchers, so that it was understood that the purpose of calibration was to position where the feedback should come on. Visibility came from the teacher's systematic approach of testing where the vibrations were and were not triggered after each calibration.

Familiar mappings. Exploit familiar ways to calibrate where possible, such as low tech materials. For example, the calibration of the Footboard and Twinkly Lights showed how the students were able to understand easily that they were setting their desired foot position by drawing around their own feet and could see how this was going to link to the functioning of the prototype. In calibration relating to the body, drawing around the body or with the body, may be a way of recording a measurement that is meaningful to the user. Familiar materials like pen and paper may not be a viable input method for the calibrating an off-the-shelf system, but it is possible to sense drawing using other methods such as touchpads. Alternatively the system could be aimed at users who want to be hands on in the process of crafting it. In which case they could draw the body parts themselves and then attach the sensors based on this.

Learning through tweaking. Learning about calibration can be done as an incremental process. This enables users to learn and understand each step of the calibration gradually before moving to the next one. Making the learning of the calibration process more incremental allows users to become confident in calibration even when the process is complex. Learning through tweaking also allows users to make a clearer link between changes they make in calibration and the outcome in terms of how the system functions.

Contextualize in the domain knowledge. This can help users set for themselves an appropriate baseline or goal using calibration. For example, the children needed to understand how to sit properly in order to calibrate the Buzzy Jackets to a good posture. This requirement for additional knowledge either needs to be designed into the calibration process so that users are informed when they calibrate or needs to accompany it in another way. The children in the Buzzy Jacket study learnt about posture from the researcher, but in another situation there could have been an instruction video which talked them through good posture while explaining how to use the system.

9.3 Summary

This chapter has outlined design guidelines that are based on the findings from the user studies in this thesis which are intended to be generalizable to the design of other real-time feedback applications. It has shown how the ways feedback is used is a complex interaction between the feedback design, the characteristics of the individual user and the setting and task it is being used for. Therefore, it is important when applying these guidelines to tailor them to the application and the setting. To do this the prototypes need to be studied in context. The next chapter discusses the methodological lessons learnt from our varying approaches to in-the-wild studies.

Chapter 10 – Reflections on Methodology

10.1 Introduction

The methodology used in this thesis has evolved as the research questions and technology have developed. This chapter reflects on the methodological outcomes learnt while conducting in-the-wild studies. The method for each study was carefully adapted to address particular research questions in a particular setting. The progression of the methodology has been towards studying real-time feedback in real-life settings and giving participants control over how they use the feedback. This is because the setting and context that participants use real-time feedback in has a considerable effect on how it is used. This lead to two important challenges: (i) how to operationalize real-life context meaningfully and practically into a study and (ii) how to understand the context that the study is taking place in, in order to understand its effect on the findings obtained.

The first challenge was addressed by using an in-the-wild methodology, in which genuine violin players were recruited to take part in the studies. The feedback was tested in a setting where they would usually practise and using their normal practice activities as motivating and useful tasks. In some cases it was not practical to conduct the study in their normal practice location - however a similar setting was used. The second challenge was addressed by an even more unusual method, namely participation in the context of the study (Chapters 5 and 7) and in some cases the study itself (Chapter 6). Participation is common in ethnography and other related methods but has very rarely been applied when testing prototypes in the wild.

While participating in the studies, I began to reflect upon my own role as researcher and how my relationship with participants affected how the study progresses, what data is collected and how it is interpreted. These reflections are summarised here. This is intended to help researchers planning a study to consider what kind of relationship they may want to build with participants and how this will affect their research. The discussion summarises the different methods that each study used and how they impacted on the results.

10.2 Summary of Methodological Evolution

The methodology used in each study was discussed at the end of each chapter and presented in the overview table below (Table 10-1).

Study (chapter)	Setting	Participants	Methods Used	Researcher Role
MuSense – comparing modalities (ch.5)	University Practice Rooms	University Orchestra Players	Within participants quantitative and qualitative study. Using contextually relevant tasks and feedback.	Encouraging Not authoritative Familiar
ShareSense – shared feedback in ensembles (ch.6)	Summer School	Adult Amateur Musicians	In-the-wild study using participant observation. Researcher participated in ensemble practice.	Explaining Encouraging Not authoritative Building familiarity
Buzzy Jacket & Twinkly Lights – programme of learning (ch.7)	High School	3rd violin/viola section of orchestra	In-the-wild study of feedback as part of a programme of learning. Children operated the feedback and took it home.	Explaining Encouraging Authority Building familiarity

Table 10-1: Overview of methods and researcher roles in the five studies conducted for the thesis

In-the-wild studies

In-the-wild studies were the primary methodological approach used to gain insight into how the practical realities of violin playing and learning impact on the use of real-time feedback (and vice versa). Accordingly prototypes were studied in the settings where they would normally learn or practise the violin. By doing this we were able to observe participants using the feedback under real-life levels of demand. These demands came from many sources, for example, playing from memory, reading unfamiliar music, working on a new goal the teacher has set or playing in harmony.

This approach was informative both from a theoretical understanding of real-time feedback and from the perspective of design. At a methodological level, the value of carrying out an in-the-wild study is that it enables both research into the interaction between a prototype and the details and demands of a particular setting. However, there are challenges to adopting an in-the-wild approach. The wildness naturally makes the outcome of a study unpredictable. Methods have to be adapted to meet the needs of the environment and this can make it difficult to produce consistent data for quantitative analysis. There are a lot of uncontrolled factors which make measurements difficult to compare. However, this does not detract from the value of conducting in-the-wild studies because they produce detailed and insightful qualitative and quantitative findings.

The studies conducted here have shown how it is possible to take key elements from an in-the-wild study while keeping some control that is normally associated with running a lab study. For example, the study with the University Orchestra ways was like a lab study because all the participants performed the same set task under a number of different conditions and measurements were taken for each. Nonetheless, the tasks and the feedback focus were contextually relevant because they had been set by the conductor. The interviews allowed insight into what it was actually like to be playing in this setting with MuSense and enabled additional contextual details to be collected. This shows that it is possible to conduct a more controlled study to collect quantitative data which still takes the context into account. However, in order to do this the researcher needs to learn about the context thoroughly before the study commences. In this case being a participant in the orchestra was valuable because I had an in depth knowledge of the context and I could observe the conductor prompting us to use more bow in rehearsals for many weeks before the study began.

In general, some form of ethnographic research of the context that the participants normally carry out the activity being studied is necessary before designing the study. This is in order to understand what is really important about that context which should be transferred into the study setting. Also it was valuable to apply the findings from our previous study of violin learning in the lab to the design of the in the wild study in chapter 5, in particular the need to make a strong link between context, task and feedback focus.

Previous Work – MusicJacket

The MusicJacket was initially studied in a lab setting and then in violin lessons with children ((van der Linden, Schoonderwaldt, Bird, & Johnson, 2011), (van der Linden, Johnson, Bird, Rogers, & Schoonderwaldt, 2011)). The findings showed that vibrotactile feedback has much potential for providing feedback to help beginners improve their initial playing technique. However, learning to play the violin requires many hours of practice over years. Moreover, one of the problems in the real world is sustaining the learner’s motivation in the context of the demands of everyday life. Every child is different in how they like to learn and practice. Teachers tailor their approach to the needs of their pupils with an appreciation for the individuality of each child. This shows the importance of considering participants as individuals rather than trying to test hypotheses that require comparing a number of participants across different conditions (i.e. control and test groups). With this in mind, the methodology adopted for the three studies reported in this thesis was to determine how real-time feedback could be used in-the-wild. I consider the outcomes of the methods used together with my role as researcher for each study in turn.

Chapter 5 – Using MuSense to Compare Modalities for Individual Practice

This study compared different modalities of real-time feedback using quantitative and qualitative data. It used an experimentally informed interpretive framework, but aimed to situate this in a naturalistic way by studying a task that was relevant to participants beyond the study setting. This study showed how it was possible to run an in-the-wild study in a controlled way – thus enabling the data to be compared in an ecologically valid setting. The most informative findings came from analysing each participant as an individual - showing how the approach to playing that each individual took affected how each used the feedback. However, this approach highlights the tension when choosing a methodology between grouping participants together to be able to measure a generalizable effect across all participants, and recognising the individuality of participants. Although, discovering a main effect which is statistically significant across all participants maybe the clearest way of showing that technology can have a positive effect on learning, this is not the only way technology might be useful to learners. Technology may enhance learning on an individual level with the effect being influenced by the experience, abilities, motivations and goals of individual participants and how demanding they find the setting they are in. To understand how technology can be useful to people in the real-world it is necessary to understand how individual differences influence the way it is used.

Researcher Role

As the main researcher in this study, I used my role as a member of the orchestra to recruit people. This meant that our relationship in the study was as friends and fellow players as well as researcher and participant. This level of familiarity meant I also empathised with the participants as players. In this role as a member of the orchestra I could not claim authority to direct them in how they should practise the orchestra pieces used in the study. However, I could use the authority of the conductor who had set long bowing as the focus for these sections to give the task credence. This shows how sometimes it is useful to discard the role of authoritative researcher, so long as other sources of expertise can be used to give the study authority.

Chapter 6 – Using ShareSense to Study Shared Feedback for Ensembles

In this study an approach taken from Ethnography was used - namely participant observation. I participated in the study using the prototypes with the other participants. I also participated in the context of the study by attending the summer school. This was effective for building rapport with participants and learning about the demands and motivation that they were experiencing in this setting. It was useful for interpreting the findings, such as understanding why the feedback worked in some situations and not others. As both designer and researcher, it gave a novel perspective as I was able to experience first-hand what it was like using one of my designs in a setting where I was under pressure. This experience shaped how I then went on to build further prototypes. This shows how the researcher's perspective on the prototypes and the study can be transformed through participation. For a researcher, participating with the other participants gives access to small but important details of what it is like to be a participant in the study and their motivations for behaving in a particular way. It also gives a shared experience which can help build rapport and a common perspective on the technology being discussed. For a designer, testing prototypes in a real life setting can make it easier to empathise with the experience of users and this can inform further design iterations.

Researcher Role

The role of the researcher changed throughout the study. Initially, it involved explaining the technology and encouraging its use. It also involved encouraging participants to take part regularly and to talk about the feedback. Gradually as the participants became more familiar with the feedback, the role of the researcher changed to one of participating as an equal within the group, making suggestions but also listening and adopting the ideas of other participants. By becoming a more equal participant, I was given an insider perspective on what it was like to use the feedback and be a member of the group, but it also made it difficult to manage. There are aspects of the researcher's role which cannot be avoided: it is necessary to collect data, ask questions about the feedback, fix technical problems and encourage them to use the feedback. These are constant reminders that although we are participating together I am not exactly the same as them. How the researcher negotiates their role in participant observation is often discussed in ethnographic work and critiques of ethnography (e.g. (Taylor (ed.), 2002)). In part it is governed by planning and reflecting on the best way to present oneself to participants, but a lot of emphasis is also placed on using the natural abilities of the researcher as a social being to interact with other people and build relationships. The position of the researcher on the periphery of the group being studied - sometimes being an insider but still remaining separate as well - while socially challenging to maintain, is also viewed as advantageous by ethnographers. This is because it gives access to the insider perspective while still allowing the researcher to take a more external analytical view as well.

Chapter 7 – Using the Buzzy Jacket and Twinkly Lights to Study a Programme of Learning

This study was the longest one conducted, where feedback was provided in two forms to high school orchestra members to help them learn how to improve their posture while practicing together at school and individually at home over a period of 14 of weeks. A primary objective was to give the children control over how they used the feedback in different settings. Several sessions were run over a period of weeks before using the feedback. This allowed us to introduce the concept of good posture gradually and through this ensure that all the children had the domain knowledge necessary to use the feedback. It also motivated them to value good posture which helped to make the feedback more relevant.

In general, when studying how new technology may be appropriated it is important not only to ensure that participants understand how to use the technology, but also that they have the relevant domain knowledge to be able to respect the value of and appropriate the technology in a way that suits them. In some domains of study, such as office work, it may be assumed that participants are familiar with tasks the technology is designed for; in other domains, such as learning, this assumption cannot be made. The approach of gradually teaching participants what they should be aiming for (good posture) before introducing the technology is one way to ensure that participants have the necessary domain knowledge to be able to appropriate the technology in an informed way. It also has the benefit of making posture a more important goal to the children than it had been previously. In some studies manipulating participants goals in this way would be considered a disadvantage; however, in this case working on improving posture was a goal endorsed by the conductor of the school orchestra and teaching the children to prioritise it could be justified as carrying out his instructions as well as helping the children to engage more fully with the technology.

Researcher Role

My role in this study was similar to that of a teacher. It involved both explaining about posture and about how to use the feedback; and encouraging, with the aim of motivating students to work on their posture and their playing in general. Later on, encouraging also took the form of questioning the children about the feedback to make them think about how they were using it. I was also in a position of authority as an adult running a session for children. This was useful for teaching aspects of posture but it needed to be managed so that the children still felt comfortable talking about how they felt. This was done by getting to know the children and building up a familiar relationship with them.

Running in-the-wild studies raises the need for reflection on role the researcher and understanding of their social contact with the participants. Recognising differences, opportunities and potential tensions helps determine how a researcher should present themselves. Factors such as age, domain knowledge, social standing, pre-existing relationships and requirements from the study such as data collection or technical support all need to be considered as part of this process. They need to be recognised and where possible used to enable the researcher to build constructive relationships with participants that lead to collecting useful data.

10.3 Context and Participation

It has been stressed how the physical setting, the kind of learning, the goals of the participants, who is involved, and their relationships with one another all play an important role in how the feedback will be attended to and used. It is argued here that these need to be accounted for and understood by choosing an appropriate methodology. This leads to two challenges:

- *How to incorporate real-world context into a user study that the researcher has limited control over*
- *How to understand the different aspects of the context of a setting*

A method used to address the second challenge was researcher participation that helped the researcher to empathise with participants.

Understanding Context – the benefits of participation

In all the studies in this thesis we found that the researcher's participation in the study played an important role in understanding context. In the first study, being an orchestra member gave insight into the different challenges that the repertoire posed and more generally the culture within the orchestra. Being an active participant at the summer school in the second study gave an understanding of the challenges and demands of that setting. Participating in the school orchestra in the last study enabled a suitable learning focus for the technology to be identified and supported with an in depth understanding of the repertoire being played and the ways the conductor wanted players to improve.

Participant observation in ethnography is used to give the researcher an insider view on the experiences and values of the people they are studying. It demands reflexivity, challenging the investigator to question their own role in the research. Current methods employed in in-the-wild studies do not normally show this level of reflexivity or consideration for the social and personal context beyond the boundary of the interaction being studied. As Dourish and Bell (2011) put it - *"how many ubicomp papers or presentations account for the author's stance?"* By joining with participants in their activities, the researcher can start to have their own perspective transformed and learn about what it is like to be a participant. For example, in Chapter 6, the researcher learnt through participation how valuable time at the summer school was to the participants and how tired they were feeling. These may seem like small findings, but this is what participation reveals – many small details. These details are invaluable as the researcher tries to interpret the actions and words of those around her. They are also important for building rapport with participants – as Rode (2011) puts it *"when in the field they [ethnographers, but, perhaps, more generally researchers] must make a hundred little decisions to gain rapport".* How could the researcher sensibly discuss how to make the prototypes more useful to participants at the summer school without this kind of contextual understanding?

Without a shared understanding of the value of time and playing together; without knowing what repertoire is being played - where the difficult parts are and what the conductor said in the last rehearsal - these things can only be known by being present and involved. These benefits do not necessarily come from taking part in the study; they come from participating in the context the study sits in, e.g. being a member of the University Orchestra in chapter 5 and attending the High School Orchestra rehearsals in chapter 7.

Participating Directly in a Study

We would argue that there is additional value in a researcher participating in the study, itself. The first-hand experience of using the prototypes under research conditions in the summer school study in Chapter 6 surprised the researcher and changed her perspective on her designs as well as her interpretation of the feedback from other participants. One may argue that such experience threatens researcher impartiality because the researcher will assume that other participants' experiences are variations of their own. However, it is common for designers and researchers to test prototypes on themselves to check that they work. These kinds of tests are often in a completely different setting, and create similar but less accurate preconceptions about what participants' experiences will be like. By using the prototypes together with participants under the same conditions, the researcher has the opportunity to both gain the first-hand experience of using the prototypes and have the generalizability of that experience directly challenged through discussion with the participants around them.

Gaining first-hand experience using prototypes, provides additional means to empathise with participants – to see the world from the participant's perspective. This understanding of another's perspective can be built between people through a dialogue but this dialogue draws on a common understanding of human experiences and the ability to imagine the self in the other's position (Wright & McCarthy, 2008). For Wright and McCarthy (2008) the value of empathy to the designer is not about becoming the user but about responding *"to what they see as the user's world from their own perspective."* Similarly, the value to the researcher of participating in a study is about gaining the breadth of experience that allows them through dialogue to imagine the world from the perspective of the participant and respond to this.

Participating in a research study needs a lot of practical consideration. A participating researcher uses a lot of their mental and physical energy in participating – they cannot expect to be as aware of things around them as they might in a purely observational study. While participating, the researcher cannot take the long-view and spot patterns as they emerge. Just as ethnographers factor time into their days to write-up notes and think over their experiences so a researcher participating in a deployment needs to allow themselves this time, ideally, with an opportunity to follow-up their thinking in the field with the same participants.

10.4 Aspects of the Researcher's Role

The researcher can play a variety of roles when conducting an in-the-wild study. However, it may be more helpful to think of the researcher's role as being characterised by their position along several dimensions, rather than discrete roles. These go beyond the idea of the researcher being either more or less involved in a study. Instead, they try to characterise the ways the researcher is involved and how this is perceived by participants. These are: (1) facilitating, encouraging or directing (2) explaining, (3) level of authority, (4) familiarity with participants, and, (5) relationship with their research. Each dimension is explained in the next section and the individual studies are placed on a scale in terms of these dimensions. This is a simplified representation as some of the dimensions, such as the researcher's relationship with research, are more complex than a single scale. However, it is still a useful way of characterising the similarities and differences between the researcher roles in each study.

Facilitating, Encouraging or Directing

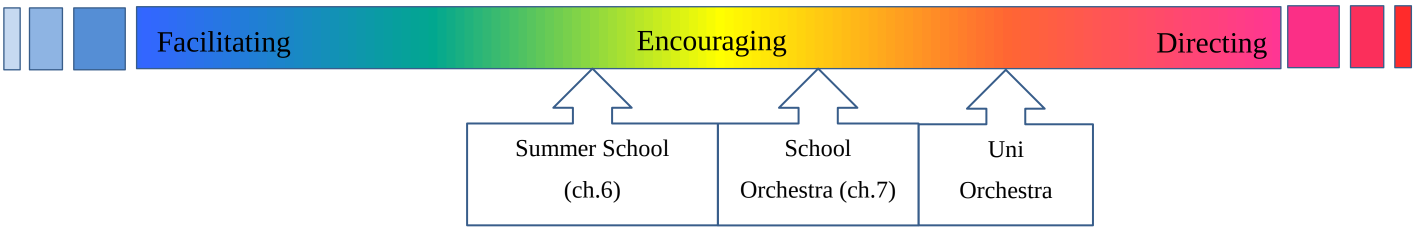

Figure 10-1: A dimension of the researcher's role is the extent to which the researcher directs the participant's actions during the study

This can be thought of as how the researcher is involved in either dictating or encouraging participants to behave in a certain way to ensure the technology is used or a study fits a certain format. At one end of the spectrum, are lab studies such as (van der Linden, Schoonderwaldt, Bird, & Johnson, 2011) where participants are directed throughout on how and when they should use the feedback and what they should play. Near to the other end of the spectrum are in-the-wild studies, such as the summer school study (Chapter 6) where participants could choose how to use the feedback while the researcher helped them with technology. In the middle of the spectrum are studies, such as the high school one described in Chapter 7, where the participants are given more control over the technology, choosing how to use the feedback but are also strongly encouraged to think about their posture and how to improve it. In other studies the researcher may be even less involved in how participants use the technology, such as (Rogers, Hazlewood, Marshall, Dalton, & Hertrich, 2010) where a technology installation was simply placed in a building and the public were free to react in any way they chose. However, even in that study one may argue that the way the technology is presented in the building is facilitating participants to interact with the building in a new way. There are also purely observational studies which do not involve the deployment of any technology; in these studies it would be possible for the researcher to consider their role as being completely hands-off.

In all the studies conducted in this thesis, the researcher did influence how participants used the technology. This is difficult to avoid when direct contact with participants since simple questions can have a subtext which implies a request. This is best illustrated with an example from my reflections about the Summer School study (Chapter 6):

"At the end of each session I asked if we could meet the next day. It is clear from the context of this question that I want to meet them tomorrow – they know I want to collect as much data as possible, it leaves them in the position where they need a reason not to come rather than a reason to come."

It is important to recognise that often a participant's motivation is to help the *researcher* as much as the *research*. As can be seen from the example above the effect of this is nuanced and the researcher must be aware of how this provides a particular subtext to all their conversations. Therefore, it is important when analysing data to think about whether certain participant behaviour is in response to suggestions from the researcher or whether it is completely self-determined. Similarly the researcher should think when making suggestions, how strongly to phrase them, and whether how to leave room for participants to disagree. This is not to argue that the researcher should minimise their involvement, because suggestions and encouragement from the researcher can act as a catalyst to enable participants to discover new ways of using the technology or to free participants to act in ways that they previously thought would not be acceptable participant behaviour. It is a valid part of the researcher's role to encourage participants to fully participate in the study, but it needs to be tempered with opportunities for participants to express their own ideas, and the influence of the researcher's encouraging role needs to be taken into account as the data is interpreted.

Explaining

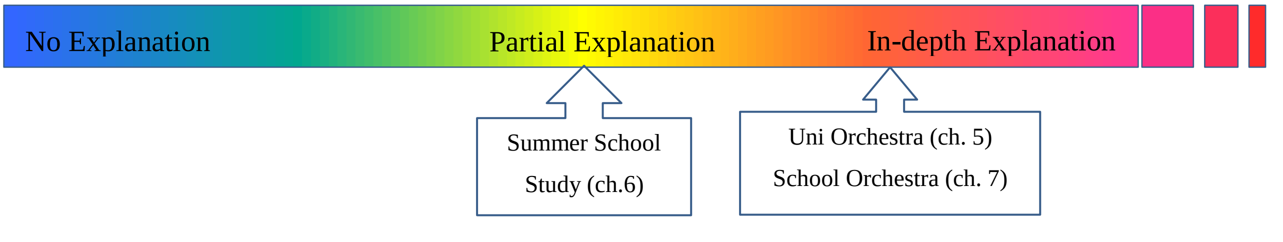

Figure 10-2: A dimension of the researcher's role is how the researcher explains the technology and the aims of the study

How much does the researcher explain how to use the system or what it is they are looking for in the study? This issue relates to facilitation and encouragement in that explaining a system facilitates its use, and depending on the context, may also encourage it too. It also implies authority as explaining how to use something shows the researcher has additional knowledge to the participant. How a prototype is explained is important because it can affect how participants use it. Within this explanation may also be nuances about how and whether participants *should* use the prototypes and this may affect how they appropriate and take ownership of them. How the study is explained is also significant because it sets up expectations in the participants about what their experience participating in the study will be like and what the researcher is looking for, these are called demand characteristics (Brown, Reeves, & Sherwood, 2011). This can then impact on the way participants behave and their perception of the study.

In all the studies it was part of the investigator's role to explain how to use the prototypes. The difficulty of this role changed with the aim of the study. For example, in the MuSense study with the university orchestra players in Chapter 5, the prototype could be described as a system to help players practise long bowing. What each display meant was described in terms of this aim, e.g. a vibration meant you had succeeded in using a long bow. Whereas in the Summer School study in Chapter 6 there was no fixed aim for the feedback even though the prototypes were very similar. The explanatory role of the researcher was a more difficult one: to give an explanation of what the prototype does without attaching a fixed interpretation to the feedback. Comparing the two studies, the goal based explanation in the university orchestra study was a more effective way of communicating, and did not prevent participants from suggesting a multitude of alternative uses in the interviews.

These insights suggest, therefore, that initially presenting prototypes through structured task-based experiences may be the most appropriate way of explaining technologies like real-time feedback. This is because having a concrete experience of using the prototype to achieve one goal appears to inspire re-imagining it in new situations; whereas an abstract explanation of what a prototype does without any concrete example of how to use it makes it hard for participants to envisage how to apply the feedback practically. In many studies giving a coherent explanation of how a prototype works is vital to the prototype being used effectively. Any explanation can be thought of as setting demand characteristics (Brown, Reeves, & Sherwood, 2011) of how participants perceive that they are expected to act and use the prototypes; however, it would be misleading to suggest that these demand characteristics are not present when no explanation is given, participants will still have preconceptions about what the technology is for and how they should behave, but these will be unknown to the researcher. Therefore in the role of explainer the researcher has the opportunity to present the technology and the role of the participant in a way which is best suited to the purpose of the study. It is important to return to this explanation again when analysing the data to consider how the explanation may have influenced participant behaviour, or whether over the course of the study participants moved beyond the researcher’s explanations to use the technology in new ways.

Level of Authority

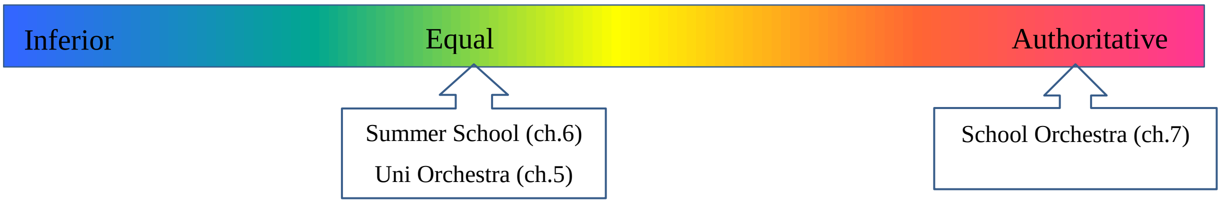

Figure 10-3: A dimension of the researcher's role is their level of authority with respect to the participant

By authority we mean the extent to which the researcher is perceived as having authority over the participant or being their equal or inferior to them, in terms of the power to direct behaviour or superior knowledge that commands respect. In traditional lab experiments, the researcher typically acts with authority over the participants, telling them what to do. The participant may also believe that the researcher has superior knowledge in the sense of knowing the true intentions of the study and making judgements about their responses. Even in a double blind study participants may perceive the researcher as having superior knowledge and authority even if this is not the case.

In the high school study in Chapter 7 the researcher was perceived as having authority because she was acting as a teacher, therefore the children considered her as knowledgeable and her arguments were given respect and credence. However, her role was different to that of a lab researcher because on the first dimension of the researcher's role she was encouraging rather than directing how the children should behave thus giving them more autonomy than a lab participant. In some cases, authority may be useful to a study; in others this inequality may obstruct building a rapport with the participants. In an in-the-wild study, authority may come from different types of knowledge and act in different ways. For example, in the study with the University Orchestra (Chapter 5) musical authority had to be 'borrowed' from the conductor to make the tasks meaningful. That the researcher was musically an equal also proved useful, as participants said they would have been more nervous playing for the conductor. However, completely relinquishing the researcher's authority also runs the risk of undermining the study; participants are giving time to the researcher and they need to believe that the researcher has the expertise to use their time appropriately. If the researcher portrays him/herself as inept or inferior in knowledge then participants may not have faith that their contribution to the study will be interpreted correctly or put to good use. Moreover it may also produce a negative attitude towards the prototypes because participants may feel the group designing them does not have the expertise to design something for a particular domain if they do not present themselves as an authority on the subject. On the other hand, it is common in ethnography for researchers to take the role of a novice or apprentice to an expert in a particular skill or culture, for example, Rosner was apprenticed to a book-binder for three months to learn both technical skills and about the bookbinder's approach to books and his craft (Rosner & Taylor, 2011). In the case of the apprentice role, the researcher is inferior in authority to the expert who is teaching them. However, the time dedicated by the researcher to building a relationship with the expert and learning from them and the novel perspective and insights the researcher may share in the process means that the expert is still likely to respect and trust the researcher.

Familiarity with Participants

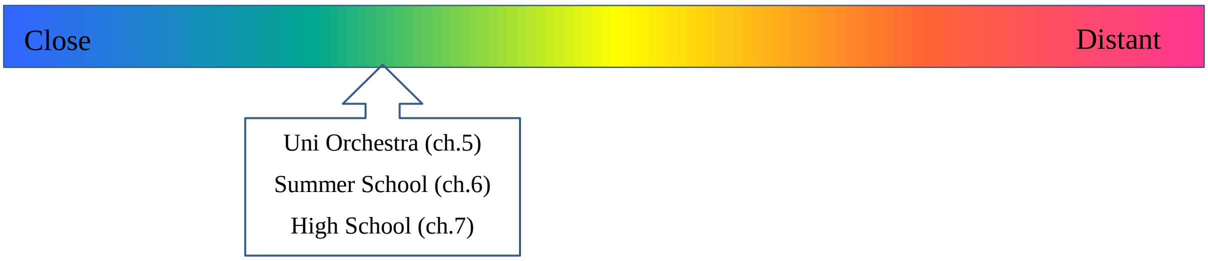

Figure 10-4: A dimension of the researcher's role is how familiar their relationship is with the participant.

In many studies, the researcher and participants are strangers to one another and never move beyond this point. In other cases, researchers may build a much closer relationship with participants or have a pre-existing relationship with them. This can potentially lead to a clearer understanding of their motivations and a familiarity with the way they communicate. It may mean that participants are willing to confide in the researcher, revealing things they would not in a first meeting. However, part of forming a friendship with participants involves being responsible towards them and the researcher must try to work out when that responsibility takes precedence over their research agenda.

In the studies reported here the researcher built a close relationship with participants. This was beneficial both personally and in terms of research. However, it does also produce anxiety when writing up, for example about analysing interactions with participants because it does not give the holistic view of the genuine goodwill from both sides. The audience that is being written for is not only the reviewers or examiners, but also the participants themselves.

Becoming familiar with the context and sharing the participants' experiences also enables the researcher to better understand and empathise with participants. For example, in the study with the University Orchestra (Chapter 5), the researcher was able to recognise and empathise with the nervousness of her participants because she was a fellow orchestra player. This enabled the researcher to react appropriately to the feelings of the participants by making jokes about her own difficulties playing the pieces to lighten the mood and show that she was not judging their playing. Gaining this kind of deeper contextual knowledge by spending time with participants and building a familiar relationship with them can facilitate better communication which can be beneficial in many ways: the researcher is able to explain the study and set up a comfortable atmosphere where participants are relaxed and act normally; the researcher can also learn how to make the study more relevant to the participants by bring in contextual details; better rapport means participants are more willing to share openly with the research in interviews and the researcher can ask more pertinent questions; and the contextual understanding provides a more valid analytical perspective when interpreting findings.

The Researcher's Relationship with the Research

Figure 10-5: A dimension of the researcher’s role is their emotional relationship with the research they are carrying out.

The researcher's feelings about the research affect the way s/he interacts with participants. For example, in Summer School study (Chapter 6), there were times where the researcher felt confident that the feedback was helpful for a participant and this made her more confident in her interactions with them. Similarly, in the High School study (Chapter 7) she felt that the programme of learning and feedback could help the children to improve their posture and this meant she was enthusiastic in the way she taught them. Participants also respond to the researcher differently depending on their perception of the researcher’s emotional relationship with the study. In the University Orchestra study (Chapter 5), the participants phrased their criticism in a way that recognised the researcher’s involvement in the design and creation of the prototypes and her future intention to improve them. They did this by suggesting improvements or balancing negative comments with positive ones.

This shows that participants are sensitive to the researcher’s relationship with the research and that having a good understanding of the researcher’s aims for the study means that they can participate in a constructive way. This comes back to the idea of demand characteristics. If the participants believe that the researcher has an emotional investment in the success of that particular prototype, they will be kind about it. If they believe instead that the researcher needs to be able to build a successful system in the future, they may be more forthcoming with criticism. This is one reason among many (Gaver, Bowers, Kerridge, Boucher, & Jarvis, 2009), (Greenberg & Buxton, 2008) why the research community needs to embrace failure as much as success; because if a prototype has to be a success for the research to be valued, then researchers will always have an emotional investment in the success of a prototype – one that they may not be unable to avoid communicating to participants.

10.5 Summary

In-the-wild user studies are important for investigating how technology and context interact. Although both elements can be understood individually using lab studies and ethnographic studies respectively, it is still not possible to predict how they will influence one another without studying technology in-situ. In-the-wild is both a challenging and rewarding methodology: challenging because it is unpredictable and requires flexibility; rewarding because it provides unexpected but important findings and a rich and interesting level of qualitative detail.

The naturalism of in-the-wild studies has a lot in common with ethnographic approaches. A key method in ethnography is participant observation. It is valued for its ability to give an insider perspective. Some of the studies in this thesis have also utilised participation for the same reason. This approach was found to be very useful for learning about context, particularly for gathering important details that would be overlooked by an outsider. It was also helpful for building rapport and empathy. We would advocate its use in future studies.

In all the studies discussed here the role of the researcher has been important. This is an often under discussed topic in research. However, it is an important one because how the researcher presents the prototype and the research will affect how participants use it and behave in the study. Moreover the rapport a researcher builds with participants affects how participants respond in interviews. We have presented five different dimensions of the researcher's role for discussion. These dimensions are designed to help those planning and analysing studies to reflect on how the researcher's interactions with participants may impact on the study and what the most appropriate role for a researcher might be in any particular setting.

Chapter 11 – Conclusions and Future Work

The main research question in this thesis was:

How can real-time feedback aid learning the violin?

Within in this question there are sub-questions concerning the design of real-time feedback, such as what modality is most appropriate to use and how much information to try and communicate to learners. There are also sub-questions about learning such as whether real-time feedback is useful in all learning settings. These questions have been answered by the findings of the empirical studies in this thesis which show that real-time feedback can help people to learn the violin, but the role that it plays is different depending on the learning context. Moreover the ideal design varies depending on the setting and the individual.

The earlier studies of the MusicJacket (van der Linden, Schoonderwaldt, Bird, & Johnson, 2011), (van der Linden, Johnson, Bird, Rogers, & Schoonderwaldt, 2011) show that real-time vibrotactile feedback can change the way people physically play the violin and that it can be incorporated into violin lessons. This thesis extends this research to show that real-time visual and vibrotactile feedback are beneficial in a variety of practice settings. The study in Chapter 5 showed that both vibrotactile and ambient visual real-time feedback can be effective in helping people to improve their playing but that there are individual differences in terms of preferred modality and the how much information to communicate. The second study in Chapter 6 showed that real-time feedback can also contribute to the learning that takes place in ensemble playing by helping players support one another in their learning. However, in this very demanding setting, feedback needs to be simple in its message and tailored to the learning objectives of the ensemble. The third study in Chapter 7 showed that real-time feedback works particularly well when it is part of a programme of learning so that it is designed to help student achieve learning objectives that they are already working on using other methods. It also shows how a multimodal system using vibrotactile and ambient visual feedback together can enable participants to work on two different elements of the same goal. These studies reveal important findings about real-time feedback and the way that different settings influence how real-time feedback can be used and what role it is most suited to.

The contribution of this thesis to knowledge can be classified in three different ways: first, to contribute to our theoretical understanding of real-time feedback in-the-wild; second, to contribute to the design of real-time feedback through practical design guidelines; third, to contribute to the methodological debate about used how to study technology in-the-wild.

The interaction between feedback characteristics and the demands of different settings as well as the individual differences between learners is explored in detail in Chapter 8. In this chapter we built a framework from the findings of the empirical studies to describe the use of real-time feedback for violin learning in terms of the stages involved in using the feedback (catching attention, taking in information, interpreting and taking action) and the factors that influence these stages which may be inherent in the design of the feedback or may come externally from the setting and the task or be individual differences in the way people perceive and respond to the feedback and the setting. This framework answers the question of how real-time feedback can aid learning the violin in a detailed and multidimensional way. It is a theoretical framework that is based in-the-wild rather than the lab and is intended to be applied to real learners negotiating real-world demands while using real-time feedback. By doing this it contributes to knowledge by giving a structured of understanding of how many factors interact to influence the way that real-time feedback helps people learn.

This research also has many practical implications for the design of real-time feedback for music learning and other related activities such as motor-learning. These were explained in Chapter 9 in the form of design guidelines. These guidelines contribute to design by taking the findings of the empirical studies in this thesis and showing how they may be applied to design new real-time feedback systems. In particular showing how real-time feedback for motor-learning in high demand settings needs to be designed differently to when it is being used as the main focus.

The third contribution to knowledge is made through discussing and reflecting on the role of the researcher when carrying-out research in-the-wild. In particular the use of participant observation, this is an accepted method in ethnography but is new to studying prototypes in the wild. This led to the reflection that many researcher roles are in some way participatory and that in-the-wild the researcher’s actions become part of the study setting. The framework for thinking about the different aspects of the researcher’s role was explained in Chapter 10.

11.1 Generalizability and Future Research

The research in this thesis focused on real-time feedback for learning the violin and other stringed instruments. However, in the discussion of the findings we produced a framework and guidelines that are applicable beyond the scope of music learning, to other motor learning applications where participants are under high cognitive load. The framework for understanding real-time feedback could potentially be applied to many other applications such as encouraging healthy posture and activity, learning sports or practicing physical therapy exercises. Further research is necessary to discover how far the factors in this framework can be applied to the in-the-wild study of other motor-learning applications, and whether new factors need to added to this model to make it more versatile.

This thesis also focused only on real-time vibrotactile and visual feedback. Therefore the aspects of the framework and the design guidelines only currently apply to these two modalities. Many of the factors and guidelines would still apply in the case of feedback in another modality but others are modality specific. Further research could be carried out to make the framework and the guidelines truly multimodal. In this case real-time feedback would need to be applied to a different motor skill where using auditory or central visual feedback is justified as they are not suitable for violin learning.

The research in this thesis has also opened up new research questions. For example, the last study began to reveal findings about the student's sources of motivation and attitudes to learning and how these affected their use of the feedback. Further research into how individual motivators interact with different designs of feedback for practice would enable a more detailed model of motivation in real-time feedback.

The most broadly generalizable contribution of this thesis is the discussion of the researcher's role within research. This is applicable to any in-the-wild study of technology and perhaps also beyond the testing of technology to the study to other types of intervention involving people such as a new training programme in a workplace or a new initiative to help people find jobs. However, it needs to be applied to more studies to be verified that it is an accurate and useful description of the researcher's role. Further dimensions of the researcher's role may become apparent in new settings or studies which demand different things from the researcher. Therefore, future research is needed to test the applicability of this model.

Bibliography

Adams, L. Gordon Training International. (2009). *Learning a new skill is easier said than done.* Retrieved October 2012, from The Business Resource: www.thebusinessresource.com

Ames, C., & Archer, J. (1988). Achievement goals in the classroom: Students' learning strategies and motivation processes. *Journal of educational psychology, 80*(3), 260.

Arduino. (2014). *Arduino Pro Mini*. Retrieved May 2014, from Arduino: http://arduino.cc/en/Main/ArduinoBoardProMini

Artist in Balance. (2012). *Violinist in Balance*. Retrieved May 2014, from Artist in Balance: http://www.artistinbalance.org/vib/

Bachour, K., Kaplan, F., & Dillenbourg, P. (2010). An interactive table for supporting participation balance in face-to-face collaborative learning. *IEEE Transactions on Learning Technologies, 3*(3), 203-213.

Bach-y-Rita, P., & Kercel, S. W. (2003). Sensory substitution and the human–machine interface. *Trends in cognitive sciences, 7*(12), 541-546.

Bach-y-Rita, P., Collins, C., Saunders, F. A., White, B., & Scadden, L. (1969). Vision Substitution by Tactile Image Projection. *Nature, 221*, 963 - 964.

Baddeley, A. (1992). Working Memory. *Science, 255*(5044), 556-559.

Baillie, S., Brewster, S., Hall, C., & O'Donnell, J. (2005). Motion space reduction in a haptic model of violin and viola bowing. *First Joint Eurohaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, (pp. 525-526). Pisa, Italy.

Bartlett, G., Stewart, J. D., Tamblyn, R., & Abrahamowicz, M. (1998). Normal Distributions of Thermal and Vibration Sensory Thresholds. *Muscle and Nerve*, 367-374.

Beilock, S. L., Carr, T. H., MacMahon, C., & Starkes, J. L. (2002). When Paying Attention Becomes Counterproductive: Impact of Divided Versus Skill-focused Attention on Novice and Experienced Performance of Sensorimotor Skills. *Journal of Experimental Psychology: Applied, 8*(1), 6-16.

Benford, S., Schädelbach, H., Koleva, B., Anastasi, R., Greenhalgh, C., Rodden, T., Green, J., Ghali, A., Pridmore, T., Gaver, B., Boucher, A., Walker, B., Pennington, S., Schmidt, A., Gellersen, H. & Steed, A. (2005). Expected, Sensed and Desired: a framework for designing sensing-based interaction. *ACM Transactions on Computer-Human Interaction, 12*(1), 3-30.

Bird, J., Marshall, P., & Rogers, Y. (2009). Low-fi skin vision: a case study in rapid prototyping a sensory substitution system. *Proceedings of the 23rd British HCI Group Annual Conference on People and Computers* (pp. 55-64). British Computer Society.

Blackwell, L. S., Trzesniewski, K. H., & Dweck, C. S. (2007). Implicit theories of intelligence predict achievement across an adolescent transition: A longitudinal study and an intervention. *Child development, 78*(1), 246-263.

Bloomfield, A., & Badler, N. I. (2008, April). Virtual training via vibrotactile arrays. *Presence: Teleoper. Virtual Environ., 17*(2), 103-120.

Boles, D. B., Bursk, J. H., Phillips, J. B., & Perdelwitz, J. R. (2007). Predicting dual-task performance with the multiple resource questionnaire. *Human Factors, 49*(1), 32-45.

Brewster, S., Chohan, F., & Brown, L. (2007). Tactile feedback for mobile interactions. *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 159-162). ACM.

Brown, B., Reeves, S., & Sherwood, S. (2011). Into the wild: challenges and opportunities for field trial methods. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 1657-1666). ACM.

Brown, L. M. (2007). *Tactons: Structured Vibrotactile Messages for Non-Visual Information Display (Doctoral Thesis).* University of Glasgow.

Buckingham, G., Carey, D. P., Colino, F. L., & Binsted, G. (2010). Gating of vibrotactile detection during visually guided bimanual reaches. *Experimental brain research, 201*(3), 411-419.

Calder, B., & Staw, B. (1975). Self-perception of intrinsic and extrinsic motivation. *Journal of personality and social psychology, 31*, 599-605.

Chapman, C., Bushnell, M., Miron, D., Duncan, G., & Lund, J. (1987). Sensory perception during movement in man. *Experimental Brain Research, 68*, 516-524.

Cholewiak, R. W., & Collins, A. A. (2003). Vibrotactile localization on the arm: Effects of place, space, and age. *Perception and Psychophysics, 65*(7), 1058-1077 .

Cholewiak, R. W., Brill, J. C., & Schwab, A. (2004). Vibrotactile localization on the abdomen: Effects of place and space. *Perception and Psychophysics, 66*(6), 970-987 .

Colavita, F. (1974). Human sensory dominance. *Perception and Psychophysics, 16*(2), 409-412.

Colavita, F. (1979). A further investigation of visual dominance. *Perception and Psychophysics, 25*(4), 345-347.

Consolvo, S., Klasnja, P., McDonald, D. W., Avrahami, D., Froehlich, J., LeGrand, L., Libby, R., Mosher, K.& Landay, J. A. (2008). Flowers or a robot army?: encouraging awareness & activity with personal, mobile displays. *Proceedings of the 10th international conference on Ubiquitous computing (UbiComp '08)* (pp. 54-63). ACM.

de Fockert, J. W., Rees, G., Frith, C. D., & Lavie, N. (2001). The Role of Working Memory in Visual Selective Attention . *Science, 291*(5509), 1803-1806.

deCharms, R. (1968). *Personal Causation: internal affective determinants of behaviour.* New York: Academic Press.

Dix, A. (2007). Designing for Appropriation. *Proceedings of the 21st British HCI Group Annual Conference on People and Computers Volume 2* (pp. 27-30). British Computer Society.

Dourish, P. (2001). *Where the Action Is: The Foundations of Embodied Interaction* (1st paperback ed.). MIT Press.

Dourish, P., & Bell, G. (2011). *Divining a digital future: Mess and mythology in ubiquitous computing.* MIT Press.

Dweck, C. S., & Leggett, E. L. (1988). A social-cognitive approach to motivation and personality. *Psychological review, 95*(2), 256.

Eccles, J. S., & Wigfield, A. (2002). Motivational beliefs, values, and goals. *Annual review of psychology, 53*(1), 109-132.

Elliott, E. S., & Dweck, C. S. (1988). Goals: an approach to motivation and achievement. *Journal of personality and social psychology, 54*(1), 5.

Ericsson, K. A., Krampe, R. T., & Tesch-Römer, C. (1993). The role of deliberate practice in the acquisition of expert performance. *Psychological review, 100*(3), 363.

Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature, 415*, 429-433.

Fleming, N. a. (2006). Learning Styles Again: VARKing up the right tree! *Educational Developments, 7*(4), 4.

Fortmann, J., Stratmann, T. C., Boll, S., Poppinga, B., & Heuten, W. (2013). Make me move at work! An ambient light display to increase physical activity. *Proceedings of the 7th International Conference on Pervasive Computing Technologies for Healthcar*, (pp. 274-277). Venice, Italy.

Frisoli, A., Salsedo, F., Bergamasco, M., Rossi, B., & Carboncini, M. C. (2009). A force-feedback exoskeleton for upper-limb rehabilitation in virtual reality. *Applied Bionics and Biomechanics, 6*(2), 115-126.

Gallace, A., & Spence, C. (2014). *In Touch with the Future: the sense of touch from cognitive neuroscience to virtual reality*. Oxford, UK: Oxford University Press.

Galvao, A. (2002). Motivation to learn: the case of professional orchestra players. *Proceedings of the 8 th Conference of The European Council for High Ability*, (pp. 62-70).

Garson, A. (1973). Suzuki and Physical Movement: Is movement the basis of all learning? *Music Educators Journal*, 34-37.

Gaver, W., Bowers, J., Kerridge, T., Boucher, A., & Jarvis, N. (2009). Anatomy of a failure: how we knew when our design went wrong, and what we learned from it. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 2213-2222). ACM.

Geake, J. (2008). Neuromythologies in education. *Educational Research, 50*(2), 123-133.

Geldard, F. A., & Sherrick, C. E. (1972). The Cutaneous "Rabbit": A Perceptual Illusion. *Science, 178*(4057), 178-179.

Geurts, L., Vanden Abeele, V., Husson, J., Windey, F., Van Overveldt, M., Annema, J. & Desmet, S. (2011). Digital games for physical therapy: fulfilling the need for calibration and adaptation. *Proceedings of the fifth international conference on Tangible, embedded, and embodied interaction* (pp. 117-124). ACM.

Glowinski, D., Coletta, P., Volpe, G., Camurri, A., Chiorri, C., & Schenone, A. (2010). Multi-scale entropy analysis of dominance in social creative activities. *Proceedings of the international conference on Multimedia* (pp. 1035-1038). ACM.

Goldstein, E. (1999). *Sensation and Perception 5th ed.*, (5th ed.). Pacific Grove, CA: Brooks/Cole Publishing Company.

Greenberg, S., & Buxton, B. (2008). Usability evaluation considered harmful (some of the time). *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 111-120). ACM.

Greene, D., & Lepper, M. R. (1974). Effects of extrinsic rewards on children's subsequent intrinsic interest. *Child development*, 1141-1145.

Grosshauser, T., & Hermann, T. (2009). Augmented haptics–an interactive feedback system for musicians. *Haptic and Audio Interaction Design* (pp. 100-108). Springer Berlin Heidelberg.

Grosshauser, T., & Hermann, T. (2009). The sonified music stand–an interactive sonification system for musicians. *Proceedings of the 6th Sound and Music Computing Conference*, (pp. 233-238). Casa da Musica, Porto, Portugal.

Großhauser, T., Bläsing, B., Spieth, C., & Hermann, T. (2012). Wearable Sensor-Based Real-Time Sonification of Motion and Foot Pressure in Dance Teaching and Training. *Journal of the Audio Engineering Society, 60*(7/8), 580-589.

Hallam, S. (2002). Musical motivation: Towards a model synthesising the research. *Music Education Research, 4*(2), 225-244.

Hallam, S., Rogers, L., & Creech, A. (2008). Gender differences in musical instrument choice. *International Journal of Music Education*, *26*, 7-19.

Haller, M., Richter, C., Brandl, P., Gross, S., Schossleitner, G., Schrempf, A., Nii, H., Sugimoto, M. & Inami, M. (2011). Finding the right way for interrupting people improving their sitting posture. In *Human-Computer Interaction–INTERACT 2011* (pp. 1-17). Springer Berlin Heidelberg.

Halonen, P. (1986). Quantitative vibration perception thresholds in healthy subjects of working age. *European journal of applied physiology and occupational physiology*, *54*(6), 647-655.

Hammersley, M., & Atkinson, P. (2007). *Ethnography: Principles in practice.* Routledge.

Hanson, J., Whitaker, D., & Heron, J. (2009). Preferential processing of tactile events under conditions of divided attention. *NeuroReport*, *20*, 1392-1396.

Hartcher-O'Brien, J., Gallace, A., Krings, B., Koppen, C., & Spence, C. (2008). When vision ‘extinguishes’ touch in neurologically-normal people: extending the Colavita visual dominance effect. *Exp Brain Res, 186*, 643-658.

Hecht, D., & Reiner, M. (2009). Sensory dominance in combinations of audio, visual and haptic stimuli. *Exp Brain Res, 193*, 307-314.

Hinrichs, U., & Carpendale, S. (2011). Gestures in the wild: studying multi-touch gesture sequences on interactive tabletop exhibits. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 3023-3032). ACM.

Hoggan, E., & Brewster, S. A. (2010). Crosstrainer: testing the use of multimodal interfaces in situ. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 333-342). ACM.

Hoggan, E., Brewster, S. A., & Johnston, J. (2008). Investigating the effectiveness of tactile feedback for mobile touchscreens. *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 1573-1582). ACM.

Holland, S., Bouwer, A. J., Dalgelish, M., & Hurtig, T. M. (2010). Feeling the Beat where it counts: Fostering Multi-limb Rhythm Skills with the Haptic Drum Kit. *Proceedings of the fourth international conference on Tangible, embedded, and embodied interaction (TEI '10)* (pp. 21-28). ACM.

Hommel, B., Müsseler, J., Aschersleben, G., & Prinz, W. (2001). The Theory of Event Coding: A framework for perception and Action planning. *Behavioural and Brain Sciences, 24*, 849-937.

Horrey, W. J., & Wickens, C. D. (2004). Focal and Ambient Visual Contributions to Driver Visual Scanning in Lane Keeping and Hazard Detection. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, *48*, pp. 2325-2329.

Huang, K., Starner, T., Do, E., Weiberg, G., Kohlsdorf, D., Ahlrichs, C. & Leibrandt, R. (2010). Mobile music touch: mobile tactile stimulation for passive learning. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)* (pp. 791-800). Atlanta, Georgia: ACM.

Huitt, W. (2001). Motivation to learn: An overview. *Educational Psychology Interactive* .

Hwang, S. H., Ara, J., Song, T., & Khang, G. (2014). A Feasibility Study to Elicit Tactile Sensations by Electrical Stimulation. *XIII Mediterranean Conference on Medical and Biological Engineering and Computing 2013* (pp. 1738-1741). Springer International Publishing.

Jambon, F., & Meillon, B. (2009). User experience evaluation in the wild. *CHI'09 Extended Abstracts on Human Factors in Computing Systems* (pp. 4069-4074). ACM.

Johnson, R., Bianchi-Berthouze, N., Rogers, Y., & van der Linden, J. (2013). Embracing calibration in body sensing: using self-tweaking to enhance ownership and performance. *Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing*, (pp. 811-820). Zurich, Switzerland.

Johnson, R., Rogers, Y., van der Linden, J., & Bianchi-Berthouze, N. (2012). Being in the thick of in-the-wild studies: the challenges and insights of researcher participation. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, (pp. 1135-1144).

Johnson, R., van der Linden, J., & Rogers, Y. (2010). MusicJacket: the efficacy of real-time vibrotactile feedback for learning to play the violin. *CHI'10 Extended Abstracts on Human Factors in Computing Systems* (pp. 3475-3480). Atlanta, GA, USA: ACM.

Johnson, R., van der Linden, J., & Rogers, Y. (2010). To Buzz or not to Buzz: Improving awareness of posture through vibrotactile feedback. *Whole Body Interaction Worshop at CHI'10.*

Kaczmarek, K. A., Webster, J. G., Bach-y-Rita, P., & Tompkins, W. J. (1991). Electrotactile and Vibrotactile Displays for Sensory Substitution Systems. *IEEE TRANSACTIONS ON BIOMEDICAL ENGINEERING, 38*(1).

Karapanos, E., Zimmerman, J., & Martens, J. (2009). User experience over time: an initial framework. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 729-738). ACM.

Kearney, S., & Fussey, I. (1991). The use of adapted leisure materials to reinforce correct head positioning in a brain-injured adult. *Brain Injury, 5*(3), 295-302.

Kempter, S. (2003). *How Muscles Learn: Teaching the Violin with the Body and Mind.* Alfred Publishing.

Kirsh, D. (2010). Thinking with the Body. *Proceedings of the 32nd Annual Conference of the Cognitive Science Society* (pp. 2864-2869). Austin, TX: Cognitive Science Society.

Konczak, J., vander Velden, H., & Jaeger, L. (2009). Learning to play the violin: motor control by freezing, not freeing degrees of freedom. *Journal of Motor Behavior, 41*(3), 243-252.

Koppen, C., & Spence, C. (2007a). Seeing the light: exploring the Colavita visual dominance effect. *Exp Brain Res, 180*, 737-754.

Koppen, C., & Spence, C. (2007b). Assessing the role of stimulus probability on the Colavita visual dominance effect. *Neuroscience Letters, 418*, 266–271.

Koppen, C., & Spence, C. (2007c). Audiovisual asynchrony modulates the Colavita visual dominance effect. *Brain Research, 1186*, 224-232.

Lake, F., & Hamdorf, J. (2004). Teaching on the run tips 5: teaching a skill. *Medical journal of Australia, 181*, 327--328.

Lamble, D., Laakso, M., & Summala, H. (1999). Detection thresholds in car following situations and peripheral vision: implications for positioning of visually demanding in-car displays. *Ergonomics, 42*(6), 807-815.

Lavie, N. (2005). Distracted and Confused: Selective attention under load. *TRENDS in cognitive science, 9*(2).

Lehmann, A. C., Sloboda, J. A., & Woody, R. H. (2007). *Psychology for Musicians: Understanding and Acquiring the Skills.* Oxford, UK: Oxford University Press.

Leibowitz, H. W., Shupert, C. L., & Post, R. B. (1983). The Two Modes of Visual Processing: Implications for Spatial Orientation. *Peripheral Vision Horizon Display.* Edwards, California: Nasa Conference Publications.

Lieberman, J., & Breazeal, C. (2007). TIKL: Development of a Wearable Vibrotactile Feedback Suit for Improved Human Motor Learning. *Trans. Rob., 23*(5), 919-926.

Lotze, M., Braun, C., Birbaumer, N., Anders, S., & Cohen, L. G. (2003). Motor learning elicited by voluntary drive. *Brain, 126*(4), 866-872.

Marshall, P., Morris, R., Rogers, Y., Kreitmayer, S., & Davies, M. (2011). Rethinking'multi-user': an in-the-wild study of how groups approach a walk-up-and-use tabletop interface. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 3033-3042). ACM.

Mayer, R. E., & Moreno, R. (1998). A Split-Attention Effect in Multimedia Learning: Evidence for Dual Processing Systems in Working Memory. *Journal of Educational Psychology, 90*(2), 312-320.

McDaniel, T. L., Goldberg, M., Bala, S., Fakhri, B., & Panchanathan, S. (2012). Vibrotactile feedback of motor performance errors for enhancing motor learning. *Proceedings of the 20th ACM international conference on Multimedia (MM '12)* (pp. 419-428). ACM.

Mulckhuyse, M., & Theeuwes, J. (2010). Unconscious attentional orienting to exogenous cues: A review of the literature. *Acta Psychologica, 134*, 299-309.

Nagel, S. K., Carl, C., Kringe, T., Märtin, R., & König, P. (2005). Beyond sensory substitution—learning the sixth sense. *Journal of neural engineering, 2*(4), R13-R26.

Neely, G., & Burström, L. (2006). Gender differences in subjective responses to hand–arm vibration. *International journal of industrial ergonomics, 36*(2), 135-140.

Ng, K. C., Weyde, T., Larkin, O., Neubarth, K., Koerselman, T., & Ong, B. (2007). 3d augmented mirror: a multimodal interface for string instrument learning and teaching with gesture support. *Proceedings of the 9th international conference on Multimodal interfaces (ICMI '07)* (pp. 339-345). ACM.

Nintendo Wii. (2006). *Wii*. Retrieved May 2014, from Nintendo: http://www.nintendo.co.uk/Wii/Wii-94559.html

Norman, D. A. (1986). Cognitive Engineering. *User centered system design*, 31-61.

O'Brien, F., & Azrin, N. (1970). Behavioral engineering: Control of posture by informational feedback. *Journal of Applied Behavior Analysis, 3*, 235.

Pashler, H., McDaniel, M., Rohrer, D., & Bjork, R. (2008). Learning styles concepts and evidence. *Psychological science in the public interest, 9*(3), 105-119.

Pielot, M., & Boll, S. (2010). Tactile Wayfinder: comparison of tactile waypoint navigation with commercial pedestrian navigation systems. *Pervasive Computing* (pp. 76-93). Springer Berlin Heidelberg.

Posner, M. I. (1980). Orienting of Attention. *Quarterly Journal of Experimental Psychology, 32*(FEB), 3-25.

Post, L. J., Zumpa, I. C., & Chapman, C. E. (1994). Perception of vibrotactile stimuli during motor activity in human subjects. *Experimental Brain Research, 100*, 107-120.

Precision Microdrives. (2014). *Pancake shaftless coin vibration motors*. Retrieved May 2014, from Precision Microdrives: http://www.precisionmicrodrives.com/vibrating-vibrator-vibration-motors/pancake-shaftless-coin-vibration-motors

Previc, F. H. (1998). The Neuropsychology of 3-D Space. *Psychological Bulletin, 124*(2), 123-164.

Read, J. C., & MacFarlane, S. (2006). Using the fun toolkit and other survey methods to gather opinions in child computer interaction. *Proceedings of the 2006 conference on Interaction design and children* (pp. 81-88). ACM.

Repp, B. H., & Penel, A. (2002). Auditory Dominance in Temporal Processing: New Evidence from Synchronization with Simultaneous Visual and Auditory Sequences. *Journal of Experimental Psychology: Human Perception and Performance, 28*(5), 1085-1099.

Rode, J. (2011). Reflexivity in digital anthropology. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 123-132). ACM.

Rogers, Y. (2011). Interaction design gone wild: striving for wild theory. *Interactions, 18*(4), 58-62.

Rogers, Y., Hazlewood, W. R., Marshall, P., Dalton, N., & Hertrich, S. (2010). Ambient influence: Can twinkly lights lure and abstract representations trigger behavioral change? *Proceedings of the 12th ACM international conference on Ubiquitous computing* (pp. 261-270). ACM.

Rosner, D. K., & Taylor, A. S. (2011). Antiquarian answers: book restoration as a resource for design. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 2665-2668). ACM.

Ryan, R. M., & Deci, E. L. (2000). Intrinsic and extrinsic motivations: Classic definitions and new directions. *Contemporary educational psychology, 25*(1), 54-67.

Salvucci, D., & Taatgen, N. (2008). Threaded Cognition: An Integrated Theory of Concurrent Multitasking. *Psychological Review, 115*(1), 101-130.

Scaife, M., & Rogers, Y. (1998). Kids as informants: Telling us what we didn't know or confirming what we knew already. In *The design of children's technology* (pp. 27-50). San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.

Schnapp, M., & Schnapp, E. (2007). Retrieved 2014, from iPosture: http://www.iposture.com/

Sharpe, J. G., Bowker, R., & Byrne, J. (2008). VAK or VAK-uous? Towards the trivialisation of learning and the death of scholarship. *Research Papers in Education, 23*(3), 293–314.

Shea, C. H., & Wulf, G. (1999). Enhancing motor learning through external-focus instructions and feedback. *Human Movement Science, 18*, 553-571.

Sichivitsa, V. O. (2004). Music motivation: A study of fourth, fifth, and sixth graders' intentions to persist in music. *Contributions to Music Education, 31*, 27-41.

Sichivitsa, V. O. (2007). The influences of parents, teachers, peers and other factors on students' motivation in music. *Research Studies in Music Education, 29*(1), 55-68.

Singh, A., Klapper, A., Jia, J., Fidalgo, A., Jimenez, A. T., Kanakam, N., Bianchi-Berthouze, N. & Williams, A. (2014). Motivating People with Chronic Pain to do Physical Activity: Opportunities for Technology Design. *CHI'14.* ACM.

Slater, J. A., Lujan, H. L., & DiCarlo, S. E. (2007). Does gender influence learning style preferences of first-year medical students? *Advances in Physiological Education, 31*, 336-342.

Sloboda, J. (2004). *Exploring the Musical Mind: Cognition, emotion, ability, function.* Oxford, UK: Oxford University Press.

Smith, B. P. (2005). Goal orientation, implicit theory of ability, and collegiate instrumental music practice. *Psychology of Music, 33*(1), 36-57.

Society of Teachers of the Alexander Technique. (2011). *What is the Alexander Technique?* Retrieved May 2014, from STAT: http://www.stat.org.uk/alexander-technique

Spelmezan, D. (2012). An investigation into the use of tactile instructions in snowboarding. *Proceedings of the 14th international conference on Human-computer interaction with mobile devices and services (MobileHCI '12)* (pp. 417-426). San Francisco, USA: ACM.

Spelmezan, D., Jacobs, M., Hilgers, A., & Borchers, J. (2009). Tactile motion instructions for physical activities. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '09)* (pp. 2243-2252). Boston, USA: ACM.

Suzuki, Y., & Kobayashi, M. (2005). Air jet driven force feedback in virtual reality. *Computer Graphics and Applications, 25*(1), 44-47.

Synertial (formerly known as Animazoo). (2014). Retrieved May 2014, from Synertial: http://www.synertial.com/

Taylor (ed.), S. (2002). *Ethnographic Research: A Reader.* Sage.

van der Linden, J., Johnson, R., Bird, J., Rogers, Y., & Schoonderwaldt, E. (2011). Buzzing to play: lessons learned from an in the wild study of real-time vibrotactile feedback. *Proceedings of the SIGCHI Conference on Human factors in Computing Systems* (pp. 533-542). ACM.

van der Linden, J., Rogers, Y., Oshodi, M., Spiers, A., McGoran, D., Cronin, R. & O’Dowd, P. (2011). Haptic Reassurance in the Pitch Black for an Immersive Theatre Experience. *Proc. UbiComp.* Beijing, China: ACM.

van der Linden, J., Schoonderwaldt, E., Bird, J., & Johnson, R. (2011). Musicjacket—combining motion capture and vibrotactile feedback to teach violin bowing. *IEEE Transactions on Instrumentation and Measurement, 60*(1), 104-113.

Van Erp, J. B., Van Veen, H. A., Jansen, C., & Dobbins, T. (2005). Waypoint navigation with a vibrotactile waist belt. *ACM Transactions on Applied Perception (TAP), 2*(2), 106-117.

Voss, M., Ingram, J. N., P., H., & Wolpert, D. M. (2006). Sensorimotor attenuation by central motor command signals in the absence of movement. *Nature, 9*(1), 26-27.

Wagemans, J., Elder, J. H., Kubovy, M., Palmer, S. E., Peterson, M. A,; Singh, M. & von der Heydt, R. (2012). A century of Gestalt psychology in visual perception: I. Perceptual grouping and figure–ground organization. *Psychological bulletin, 138*(6), 1172-1217.

Weiner, B. (1985). An attributional theory of achievement motivation and emotion. *Psychological review, 92*(4), 548.

Weinstein, S. (1968). Intensive and extensive aspects of tactile sensitivity as a function of body part, sex and laterality. *First Int'l symp. on the Skin Senses*, (pp. 195-218).

Weiser, M. (1991). The computer for the 21st century. *Scientific American, 265*(3), 94-104.

Weiser, M., & Brown, J. S. (1996). Designing calm technology. *PowerGrid Journal, 1*(1), 75-85.

Weiss, S. M., Reber, A. S., & Owen, D. R. (2008). Locus of Focus: The effect of switching from a preferred to a non-preferred focus of attention. *Journal of Sports Sciences, 26*(10), 1049-1057.

Wickens, C. (1980). The structure of attentional resources. In R. Nickerson (Ed.), *Attenbtion and Performance VIII* (pp. 239-257).

Wickens, C. (2008). Multiple Resources and Mental Workload. *Human Factors: The Journal of the Human Factors and Ergonomics Society, 50*(3), 449-455.

Wickens, C. D. (2002). Multiple resources and performance prediction. *Theoretical Issues in Ergonomic Science, 3*(2), 159-177.

Williams, J. M., & Andersen, M. B. (1997). Psychosocial influences on central and peripheral vision and reaction time during demanding tasks. *Behavioral Medicine, 22*(4), 160-167.

Wright, P., & McCarthy, J. (2008). Empathy and experience in HCI. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 637-646). ACM.

Wulf, G. (2007). Attentional focus and motor learning: A review of 10 years of research (Target article). *E-Journal Bewegung und Training*, 1-11.

XBox Kinect. (2010). *Kinect*. Retrieved May 2014, from XBox: http://www.xbox.com/en-GB/Kinect

Xiao, X., & Ishii, H. (2011). MirrorFugue: communicating hand gesture in remote piano collaboration. *Proceedings of the fifth international conference on Tangible, embedded, and embodied interaction (TEI '11)* (pp. 13-20). ACM.

Zdzinski, S. F. (1996). Parental involvement, selected student attributes, and learning outcomes in instrumental music. *Journal of Research in Music Education, 44*(1), 34-48.

Zheng, Y., & Morrell, J. B. (2013). Comparison of Visual and Vibrotactile Feedback Methods for Seated Posture Guidance. *IEEE Trans. Haptics, 6*(1), 13-23.

Appendix A – Designing the MusicJacket

A.1 Introduction

To design real-time feedback for learning the violin it is necessary to understand what is involved in playing and learning the violin. This Appendix describes the preliminary research carried out to investigate violin learning to find where real-time feedback could be applied to make learning easier. We began by conducting interviews and observations with violin teachers and reading around the subject on forums and in teaching materials. Our findings from this are a description of how the violin is played and common mistakes which learners make. Using this we designed the MusicJacket, a jacket which uses vibrotactile feedback to tell learners about their bowing. This design was then evaluated through expert consultation and in a small scale user study to improve the design.

A.2 Preliminary Research

Preliminary research was conducted investigating playing the violin and how it is taught. It aimed to answer the three research questions:

1) What is involved in playing the violin?

2) What aspects of playing violin could real-time feedback be applied to?

3) What modalities could be used and what are their advantages and disadvantages?

A.2.1 Method

A variety of methods were used to gather data about playing the violin. Firstly I had fourteen years' experience of learning and playing the violin and reflecting on this gave me a basic understanding of the actions and techniques involved in playing and some knowledge of bad habits and mistakes which can be made while playing. Two violin teachers were interviewed and their lessons were observed. Both teachers mainly taught children, but also taught adults as well. Teacher A taught children from as young as 6 years old, and the class observed was this age group. The lesson observed with teacher B was with a ten year old. The interviews give a teacher's perspective on playing the violin. They explained the ideal way to play the violin, and, based on their many years of experience, they described common mistake that were made by learners and the ways they would try to fix these. Observations of lessons also enabled us to see their teaching techniques first hand and understand better the challenges faced by children in the early stages of learning the violin.

Data was also collected by looking at online resources about learning and teaching the violin, in particular violinist.com discussion boards were interesting for discovering and examining debates about how the violin should be played as there is no universally agreed method. Books for violin learners and published discussions about violin teaching methods were also used to gain another perspective about violin teaching and learning, these have already been discussed in the literature review. All these sources of data about playing and teaching were combined together to form a description of the main points of playing the violin and common mistakes that can be made when learning.

A.2.2 Terminology

To explain playing the violin it is necessary to know the words used to describe the different parts of the violin. Figure A-1 shows the parts of the violin that will be referred to in the rest of this thesis. When bowing the violin a bow stroke from the heel to the tip is known as a *down bow* and a bow stroke from tip to heel is known as an *up bow*.

A.2.3 Playing the violin

The violin is supported by the collar bone, the left shoulder and the left hand (Figure A-4). Some methods prescribe that the violin should be entirely supported using the collar bone, shoulder and chin to grip it, whereas other methods say that this produces too much tension in the shoulder and jaw particularly for beginners. Violin teacher B taught her students to hold the violin mainly with the shoulder and chin, whereas teacher A taught her student to support it with the left hand. Both teachers encouraged their pupils to use a shoulder rest or sponge on the underside of the violin to help support the violin on their shoulder. However, some people reject the use of shoulder rests, as can be found in the frequent discussions of this topic on violinist.com. The neck of the violin rests on the thumb and fingers which are curled over the fingerboard and used to play different notes on the string. The wrist and hand should be in line with the lower arm. It is often tempting when starting out, to try and support the violin using the wrist and base of the hand. However, this position will prove troublesome as learners try to progress to more complicated pieces. To discourage this one of our teachers told her younger students to imagine a mouse running up and down their arm, if the student brought their wrist up to the neck of the violin then it would squash the mouse. There is variation in different methods about how the violin should be held. Teacher A taught her pupils to hold the violin so that it points forward with the elbow over the left foot which is slightly forward, teacher B taught her pupils point the violin out towards the side. Both teachers agreed that the left elbow should hang loosely underneath the violin, pointing down towards the ground. The scroll which is the carved wood on the end of the violin should point out horizontally as shown in Figure

A-2. When people start to learn the violin their left arm may become tired and let the violin slide down so that the scroll points towards the ground (Figure A-3). However, this should be avoided as it makes it more difficult to bow and play in tune. This was a particular difficulty for the pupils we observed with teacher A, because they were so young. However, I also struggle with this problem even though I have played for many years.

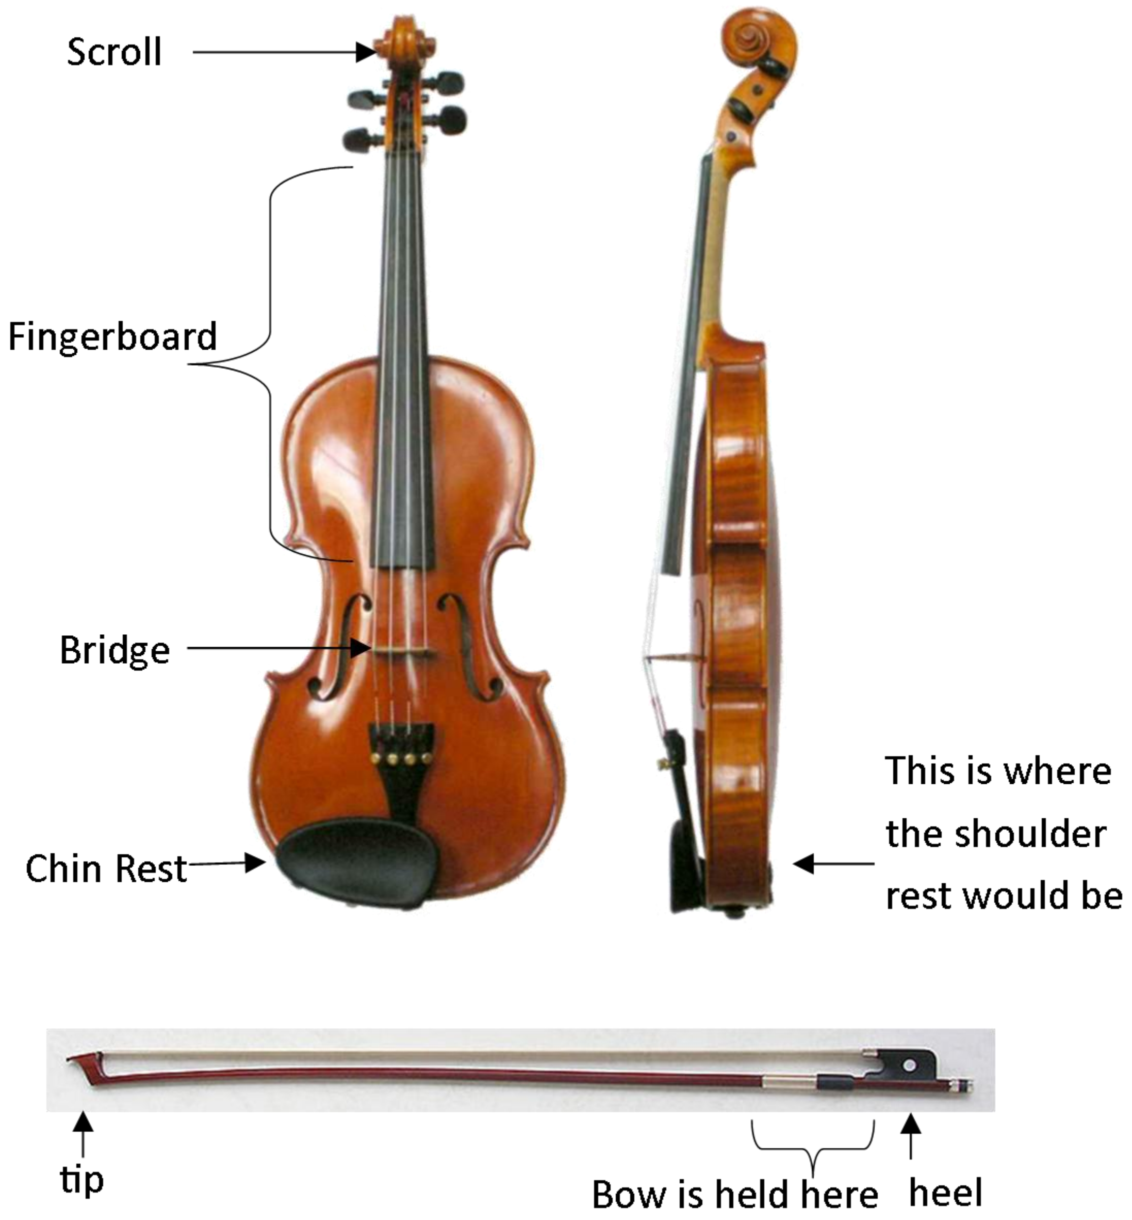

Figure A-1: Parts of the violin. Picture of the violin taken by just plain Bill and used under creative commons license CC0 1.0. Picture of the bow taken by Georg Feitscher and used under creative commons licence Attribution-ShareAlike 3.0

Different notes are played by pressing the strings down on the fingerboard with the fingers of the left hand. The pitch of the note is based on the position of the finger on the fingerboard. Unlike the guitar, the violin does not have frets, so the player must make sure that their fingers are placed in the exact position to play the notes in tune. Players must listen carefully to make sure that they are playing in tune. One of the reasons Teacher B told her students to hold the violin with the scroll pointing out to the side was that this brought the violin closer to the ear so that the pupil could be more connected with the sound. Initially most learners begin playing in first position, this is where the left hand is positioned near to the scroll end of the neck of the violin and the fingers are used to play the lower notes on the strings (this is the hand position in Figure A-4). This approach is the case in most violin teaching materials and was also used by the teachers we observed[1]. As a player progresses they will need to move their hand into other positions further up the fingerboard of the violin towards the body. This is why it is so important to hold the violin correctly because this will facilitate quick movement between different positions.

Holding the bow is a complicated position and was something that both teachers spoke about in their interview. The bow is held loosely in the right hand. There are a number of different bow holds, but the most widely used for beginners is the Galamian bow hold. The bow is held in the joints of the fingers, with all fingers bent and the little finger tip placed against the bow (see Figure A-4). The main aim of a bow hold is to enable maximum flexibility and control with minimum tension in the hand. To play the violin, the bow should be positioned between the bridge and the fingerboard and move perpendicular to the string (Figure A-4). This is a difficult movement and teacher B worked on it with her student during the observation. She had to resort to putting her finger on her pupil's violin to stop her from sliding her bow onto the fingerboard. If the bow moves from the desired position on the string then the sound will change and potentially get squeaky. To make the correct bowing movement the player moves a little from the shoulder first and then extends from the elbow (Figure A-5). Often beginners only move their arm from the shoulder joint and do not extend from the elbow, resulting in crooked bowing. Teacher B demonstrated both correct and crooked bowing to help us to understand the key differences (Figure A-6).

Playing the violin is a complex motor skill where subtle differences in posture and movement affect the sound produced. For example the position of the bow between the fingerboard and bridge can be used to make different effects: playing close to the bridge makes a raw loud sound, playing close to the fingerboard makes a dull wispy sound. Volume is controlled by bow speed and by tilting the bow so that more or less horse hair is in contact with the string. Articulation can be added by lifting the bow off the string so that it strikes the string at the start of a note. Playing with different parts of the bow also affects the sound. Playing near the heel allows the player to strike the string with more force and control and to play louder. Playing

[1] There are exceptions to this method see http://www.violinist.com/discussion/response.cfm?ID=19741

near the tip is generally quieter and smoother. Bowing alone takes approximately 700 practice hours to learn the basics skills involved (Konczak, vander Velden, & Jaeger, 2009).

Figure A-2: Correct hold of the violin with scroll pointing out horizontally

Figure A-3: Incorrect violin hold, with scroll pointing downwards

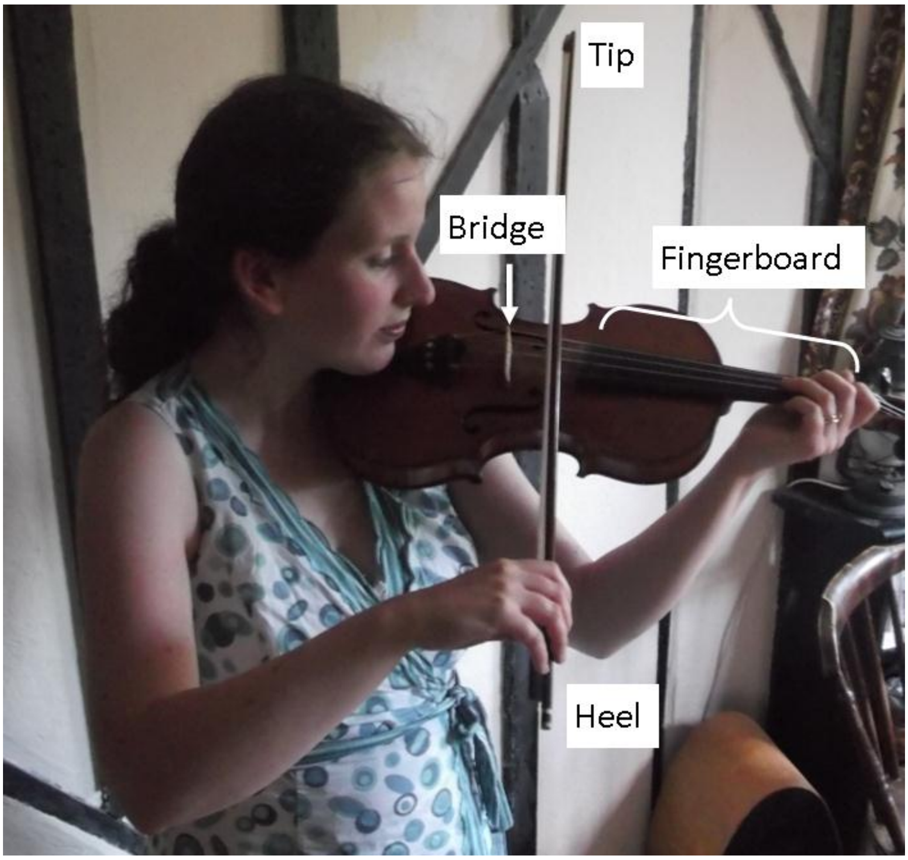

Figure A-4: When bowing the violin the bow should move perpendicular to the string

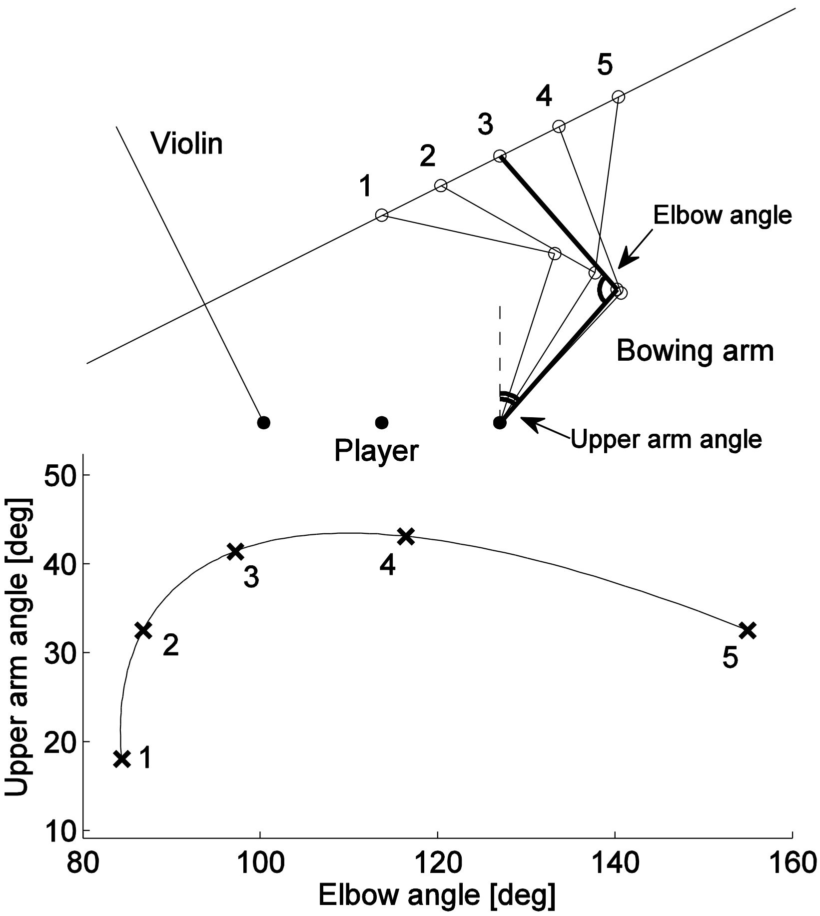

Figure A-5: Angles of the arm when playing a straight bow. The x-axis shows the movement of the players elbow joint, the y-axis shows the movement of the shoulder joint. To bow in a straight line the player must first move from their shoulder joint and then extend

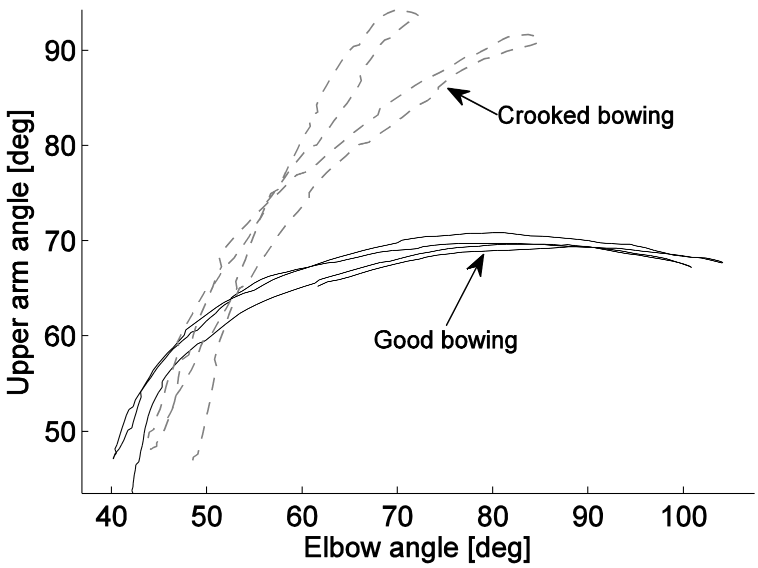

Figure A-6: Examples of the arm angles used in good bowing and crooked bowing as modelled by a teacher. In her good bowing she is moving her elbow joint more than her shoulder. In the crooked bowing she is mainly moving her arm using her shoulder joint resulting in rounded bowing.

A.2.4 Possible Aspects of Violin Playing for Feedback

The teachers highlighted many difficulties that students face when learning violin. Issues with bowing came up commonly and in one of the observations the teacher spent much of her time trying to make the student more aware of her bow and teach her to bow straighter. In another lesson with younger students the focus was more strongly on violin hold, in particular remembering to hold it up and keep it under the chin. Table A-1 gives a list of different areas aspects of technique feedback could be given about. Some like straight bowing and holding the violin concern gross motor skills whereas others like fingering notes are much more fine-grained. This distinction will be important when it comes to considering what it is possible to sense without encumbering the player and how information should be delivered to players.

Learning Goal	**Description**	**Skills involved**
Straight Bowing	Keeping bow in correct position between the bridge and fingerboard. Keeping bow perpendicular to strings.	Gross Motor
Amount of bow	Players need to be able to play long bows using the full extent of the bow and short bows using only a small amount of the bow. One of teacher B's pupil found full bows challenging.	Gross Motor
Holding the violin	Players need to hold the violin up and hold their arm in the correct position.	Gross Motor
General Posture	Players need to stand or sit up straight with a secure footing and minimum tension.	Gross Motor
Playing Together	Players also need to learn to play together, synchronising their bowing and keeping in time with one another.	Vision, Listening, Gross Motor
Fingering notes (tuning)	Learning the correct positions on the fingerboard to play different notes is very important. However the violin already gives real-time feedback about this which players needs to learn to listen to.	Listening, Fine Motor
Bow Hold	Holding the bow in the joints of the fingers with all fingers bent and hand rounded.	Fine Motor

Table A-1: Challenging aspects of violin playing which feedback could be given about

A.2.5 Constraints on Modality of Feedback

Playing the violin already has a high demand upon the senses and the violin already gives real-time feedback to the player in the form of sound and vibration. The primary way a student monitors their own playing is by listening to the sound they are making. This is particularly important for checking the tuning of notes. Students will also be able to feel the vibration of the string and the violin through their left hand and collar bone and their bow in their right hand. Students may not be so aware of this and the teachers did not mention it in their lessons but it may be important. Players need to be able use to their visual focus to read music. The younger children were not very reliant on the music as they were only beginning to learn to read it, but more advanced students like the one we observed with teacher B are reliant upon following the written music.

This places constraints upon the modality which could be used to give feedback. As listening to the violin is so essential, it would not be appropriate to give auditory feedback which might interrupt the player's connection with the sound they are making. Visual feedback is a possibility but it would have to use the peripheral vision as the main visual focus is already taken. Tactile feedback is a promising option but there is a chance that the vibrations from the instrument may make tactile feedback harder to monitor.

A.2.6 Design Criteria

From this preliminary research key design criteria governing how a feedback system for novice violinists should be designed were identified:

- Customisability: The research revealed that there is no consensus on some details of how the violin should be played. Therefore teachers should be able to calibrate the system to their preferred playing method.
- Portability: Both teachers travelled to their students, therefore a system needs to be portable and relatively quick to set-up.
- Modality: Feedback should not prevent learners from connecting with the feedback from the violin or reading music therefore feedback should either be tactile or peripheral visual feedback.
- Freedom of Movement: As violin playing involve large movements and even subtle changes can affect the sound it is important that the system should not encumber the player.

A.3 Design

Straight bowing and violin hold were selected to be the focus for real-time feedback in the initial design. Straight bowing is something that came up particularly in interviews and observations with teacher B and is an important and challenging skill that needs to be mastered to progress with playing the violin. Violin hold came up as important in the interviews and observations with teacher A and is a basic but necessary skill that all violinist need to learn.

To be able to give feedback about playing it is essential that the system is able to sense the movement of the player and judge whether they are playing correctly or making mistakes. Straight bowing and violin hold both involve gross motor skills which means that they can be sensed using a wearable sensor system without encumbering movement. In contrast in a fine motor skill like fingering notes on the violin could not easily be sensed with a wearable system as this would hamper the finger movements needed to play correctly. A wearable system was chosen rather than a vision based system in order to make the system portable. Most vision-based motion tracking systems take a long time to set-up and calibrate. In addition the violin would occlude parts of the body which would hinder tracking. Sensing could have been have also been achieved by instrumenting the violin itself, however we did not want to do this as it would change the weight and balance point of the violin and bow.

Due to the demands of playing violin either peripheral visual feedback or vibrotactile feedback are the most suitable modalities for giving information about playing. Tactile feedback was chosen for this prototype as it appeared to have the most potential for teaching physical skills. As already discussed in the literature review, tactile feedback has been found to be more effective than visual feedback for motor learning. Moreover tactile feedback can be given on the body and can draw attention to the part of the body which is making the mistake, which should make it easier to interpret than visual feedback which separate from the body.

Tactile feedback can be given in different ways: vibrations using coin vibrators (e.g. (Bird, Marshall, & Rogers, 2009)), taps on the skin using a solenoid, heating or cooling using the Peltier junction, or even tactile sensations from electrical pulses (e.g. (Hwang, Ara, Song, & Khang, 2014)). Vibrators were selected because they give continuous feedback rather than discrete taps like a solenoid would. Vibrators give more urgent and immediate feedback than temperature which takes time to change and needs to be in close contact with the skin to give an immediate sensation. Vibrators are less intrusive than the electrical stimulation which could potentially deter participants from taking part especially since one of the desired user groups is children. Vibrators have the drawback that they can confuse the player's link with the vibration of the string. It was felt however that this link was not as important to novices in their early years of playing as high level players. The feedback was set up so that it pointed out when the pupil was making a mistake; meaning when they were playing well they would receive no feedback. This setup was adopted because the vibrations were considered mildly irritating, something that would be desirable to avoid rather than aim for. Having made these design decisions, an initial prototype called the MusicJacket was built that used wearable motion capture technology to sense the movement of the upper body and coin vibrators to give feedback to players when they made a mistake in their bowing violin hold.

A.3.1 MusicJacket - First Iteration

The initial prototype is shown in Figure A-7[2]. The motion of the player is captured using the Animazoo™ inertial motion capture suit and software. The suit measures motion through accelerometers, gyroscopes and magnetometers and outputs the position of the limbs of the

[2] Please note: This is not only my work. The MusicJacket was designed, built and tested by a small group of researchers: Erwin Schoonderwaldt, Jon Bird, myself and my supervisors.

upper body in relation to the small of the back by measuring the angles between different parts of the body and mapping these onto a skeleton model. The position of the violin was inferred through looking at the position of the left shoulder and hand. The trajectory of the right hand is used as the trajectory of the bow. The movement of the player is analysed at approximately 25 frames per second and the deviation from the ideal violin position and bow trajectory is calculated. The ideal position comes from a calibration (see next section). If the player is outside a certain tolerance of the ideal position then the vibrators on the arms are switched on to guide the hand back to the correct place. It is intended to work using the metaphor that the vibrators are pushing the hand or arm.

If the pupil is playing in 1st position on the violin, the position of their left hand has two degrees of freedom: vertically and horizontally side to side. Similarly the right hand can deviate from the ideal path along two axes: vertically and horizontally back and forth. Each degree of freedom was given two vibrators, one to push the hand one way along the axis and another to push it the other. For example, if the pupil let their violin droop down they would receive a vibration underneath their arm and if they held it too high they would receive one on the top of their arm. The position of the vibrators is shown in Figure A-7 with corresponding instructions underneath. The vibrations were delivered by 10mm dc shaftless motors controlled by an Arduino microcontroller connected to a laptop via a USB. They felt very similar to those given by a mobile phone. The motion capture suit also sent its data to the laptop wirelessly and a program written in open frameworks used this data to decide when to activate the motors.

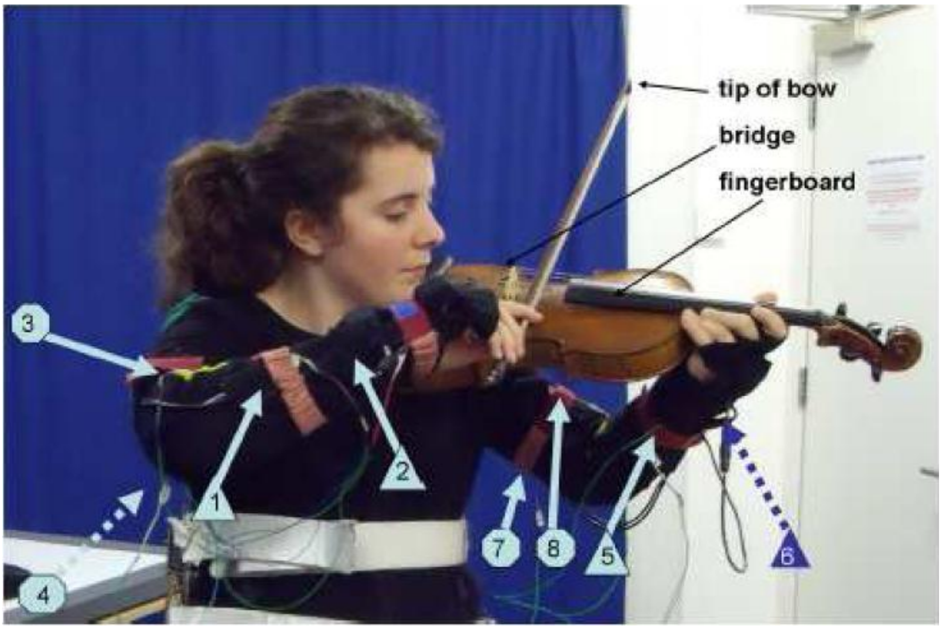

Figure A-7: First prototype (1 move bow hand forward, 2 move bow hand back, 3 move bow hand down, 4 move bow hand up, 5 move violin left, 6 move violin right, 7 move violin up, 8 move violin down)

Calibration

It was important to have a calibration process for two reasons. First, the motion capture technology needs to be calibrated to get accurate measurements. Second, calibration can also allow the teacher to set the MusicJacket to teach their method of playing the violin. Before using the suit each participant had to be photographed standing in a cube of specific dimensions in order to calibrate for their individual limb lengths. When the suit is switched on it must be kept stationary for several minutes while the inertial motion sensors calibrate for drift. Once the jacket has been put on, the wearer must face north standing up straight for a moment so that the sensors can orientate themselves.

Next an ideal position for the violin is recorded by the MusicJacket as the pupil holds the violin under guidance of the teacher. Feedback can then be given based upon how close their hand is to the correct position. After this, a reference bow stroke (Figure A-8) is recorded to show the system what the ideal playing movement would be. This was done in two ways. The teacher would either guide the pupil's hand in several bow strokes, or the teacher would hold the bow on the string and the pupil would run their hand up and down it. A line is fitted to this path in a co-ordinate system in which the violin is one of the axes. This means that if the violin is moved, the correct path for the bow will follow.

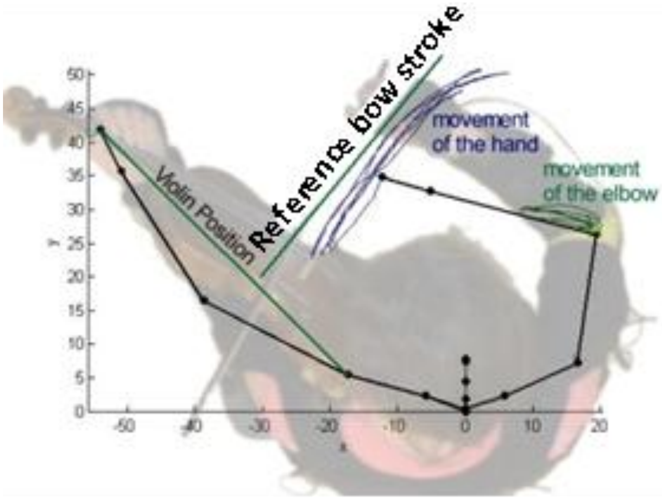

Figure A-8: Bow stroke as performed by a player overlaid with the reference bow stroke

A.3.2 Iterating on the Design

The prototype design was evaluated by (i) a small scale user study and (ii) a demonstration with an experienced music teacher and a specialist in the Alexander Technique.

(i) Initial User Study

Four volunteers took part in the user study. They had never played the violin before. Prior to using the prototype they each received a 40 minute lesson. This consisted of showing them how to hold the bow and violin and some simple rhythmic exercises that used different parts of the bow. After this a baseline recording of their motion playing each exercise was taken. The participants then had the opportunity to practice their exercises with the MusicJacket for approximately one hour spread over two or three sessions. After this another recording was made of their playing motion and for some a third a few days later. Most of the sessions were videoed and their reactions to the feedback were observed during the sessions as well.

The qualitative results from this study, described in the next section, were used to improve the design of the MusicJacket. Quantitative data about the participants bowing movements and violin hold was also collected using the motion capture system. However due to technical faults this data was incomplete for most of the participants. One participant produced a full set of data; this has been aggregated with data from a later user study (Appendix B).

Qualitative results

The participants found it overwhelming to have vibrotactile feedback about both violin hold and bowing at the same time. So for the majority of the study we chose to concentrate on feedback about bowing as this seemed to be more challenging for these adult participants. The learning experiences of each participant were found to be quite different, depending on their 'profile' and are presented individually below. The key findings are then summarised afterwards.

Participant A found it difficult to achieve straight bowing during the initial lessons. She has a petite build which meant she needed to extend her elbow further than feels natural as she reached the tip of the bow to keep the bow straight. However, she was always inclined to move from the shoulder and keep her elbow bent which causes the bow to move in a curve rather than a straight line. She also wanted it to be noted that she is a "very tense person" which would contribute to the elbow remaining bent. The effect of the feedback on A was substantial. Within a few minutes of receiving the feedback she began to extend her elbow much further. As well as this she began to evaluate her own weaknesses using the feedback, for example, she saw that it was always at the tip of the bow where she had most difficulty keeping the bow straight.

In the second feedback session A remembered clearly the need to extend the elbow and therefore began the session playing much straighter than she had at the start of the first session. When she received vibrations on the lower arm she tried to extend her arm further and was confused when this did not stop the vibrations. In fact the vibrations were actually indicating that she needed to move her hand back and extend her arm less. Her confusion could have been partly due to the vibrator being repositioned between the two sessions to aid comfort. Once she realised her mistake she found it easy to find the trajectory where she felt no feedback. “I'm fixing it immediately now” she pointed out with satisfaction. She described the vibration as “annoying and you want to make it stop but it definitely works” Towards the end of the first session she commented on how her lower arm felt “itchy” and “weird”.

Participant B In the initial lesson B learned to hold and bow the violin quickly. He looked at where his bow met the strings and monitored this as he bowed resulting in very straight bow strokes from the beginning. In his first feedback session B received less exacting feedback than A due to a mistake in the calibration of the motion capture suit. This meant he could deviate further from the ideal trajectory without getting feedback. During this session he reported receiving feedback at the start of exercises then as he adapted himself to a straighter path the feedback would stop and would not return for the rest of the exercise. He said he did not feel conscious of the opposing pairs of vibrators pushing him to the correct place, more that getting feedback “focuses attention that it's not right” and that he already had an idea of where it ought to be. He felt that having the vibrotactile feedback meant that he could shift his gaze from the bow and look around more which would be a useful outcome for learners playing from sheet music.

In the second feedback session B received more accurate and exacting feedback due to a better calibration. When he bowed on the lower half of the bow he was over extending his arm, but once the feedback was introduced he quickly brought his hand back to the correct position. When he bowed near the tip, the opposite was true, and once the feedback was triggered he brought his hand forward. In one of the exercises he also began to talk more about being “pushed” in either direction by the feedback as he over compensated, showing the concept of opposable pairs was beginning to come into play. When the MusicJacket was switched off there was a tendency for bowing to become more rounded again. Similar to A, when asked about the sensation of the vibrations he described them as “annoying - there is a strong motivation to make them stop.”

Participant C Part of C’s motivation for taking part in the study was the chance to learn to play the violin. In the lessons his bowing was inconsistent because he was experimenting to see how the sound changed as he bowed differently. Throughout the sessions his arms were tense and his bowing seemed quite uncontrolled as a result. He found it difficult to know how to react to the feedback and was not always certain which vibrator was going off. He felt he would like some time away from the violin to learn how to react to the vibrations because “it's not intuitive”. Towards the end of the first session he did not feel that he was giving the feedback any attention at all because he was concentrating on “getting the bow to work to make the sound”.

In the second session the feedback delivered was unhelpful due to a number of technical problems with the prototype. C commented on how the vibrations were not giving him sensible information. So although he had begun to understand the feedback much more, he lost faith in the system because it led him onto a bad trajectory. In both sessions he was most interested in the sound of the violin and it was with this that he monitored his success and motivated himself. This may be because he comes from a very musical background. In contrast to A and B, C never complained about the sensation of the vibrators as being uncomfortable. At the end of the second session he happened to see a recording of himself playing on screen from the motion capture system and suggested this might be more useful to see this in real time whilst learning rather than having the vibrotactile feedback.

We observed that Participant *D* learned quickly in the lesson sessions. Before the feedback commenced she could already bow with quite a straight trajectory when she was focused. She received three shorter feedback sessions because she found playing tiring. In the first feedback session, she found it difficult to understand what the vibrations were telling her and it took conscious effort to work it out. In the second session she began to understand better and was able to practise all the exercises with vibrotactile feedback. By the end of the second session her posture holding the violin was lopsided and looked uncomfortable. This may have been caused by her reactions to the vibrotactile feedback.

In her last feedback session, D understood the feedback much better and her bowing began to improve. She began to extend her arm slightly more towards the tip. When she found the trajectory which gave no feedback a sense of achievement could be seen in her facial expression. She saw the lack of vibrations as a measure of how well she was playing. After managing to play with no feedback she exclaimed slightly jokingly “I was doing it perfectly!” Once the vibrotactile feedback was removed she had a tendency to revert back to rounded bowing. D described the vibrations as “tickly”. She was more concerned about the discomfort caused by fatigue from playing than the vibrations.

In summary, the following are key findings to improve the design of the MusicJacket:

- Participants did not find the position of the vibrators intuitive to understand the message they were intended to communicate.
- Learning how to respond to the feedback affected how satisfying and effective the sessions were for participants.
- Being given information about two things at once (e.g. bowing and posture) by the feedback is confusing.
- Three of the participants displayed a sense of achievement when able to play without triggering the feedback.

- For one participant the feedback really improved her playing. This seemed to be because straight bowing was her particular challenge.
- Another participant found the feedback difficult to understand and restrictive because it did not allow him to experiment with the way he played.

(ii) Expert Consultation

The prototype was tried out by an experience music teacher and a violinist who specialised in Alexander Technique (Artist in Balance, 2012). This was to get feedback about the design and recommendations for improvements from people experienced in teaching and playing the violin. These experts were able to give a different perspective on using the MusicJacket that our user study participants, because they had greater knowledge about what is involved in learning the violin and the challenges that players face as they progress. The music teacher echoed the need for the wearer to have control over when they receive feedback. She wanted pupils to be able to switch the feedback off so that they could rest their arms without feedback when they wished. She also felt that the prototype was too strict particularly with the bow arm and that variable error margins would be useful so that the suit could be customised for different levels of user.

The Alexander technique "teaches the skilful 'use of the self'" (Society of Teachers of the Alexander Technique, 2011) mainly concentrating on the relationship between the head, neck and back. It is popular with musicians suffering from musculoskeletal problems related to playing. Due to her training, the Alexander technique specialist was acutely aware of her body which made her well suited to evaluating the effect of the vibrators. As an expert violinist and teacher she also knew the common mistakes players make. On trying on the first prototype her reaction was that it made her feel too "army". By this she meant that due to the positioning of the vibrators it focused her attention too strongly on her arms drawing attention away from her core (the back, neck and head). She also noted that positioning the vibrator on the top of bicep caused her muscle to contract working against the intended push downward. A variety of new positions for the vibrators were tested. She described what she felt and her instinctive reaction to the vibrations in that position.

Together we determined a new set of position for the vibrators (see Figure A-9). It was found that very subtle differences in position could produce very different reactions to the vibrations. For example when the vibrator was placed either side of the torso just below the ribs she found it made her diaphragm vibrate which was an uncomfortable sensation; however, moving the vibrators a few cms up to lower ribs produced a feeling of "lift". To summarise, the consultation with experts led to the following findings that can be directly applied to improving the design:

- Pupils need to be able to switch the feedback off themselves so that they can rest.
- Feedback focuses too heavily on the arms drawing attention away from core posture.

- Vibrators placed on muscles like the bicep makes them want to contract and are uncomfortable.
- It is difficult to distinguish between two vibrators placed on the same bone.

A.3.3 Second Iteration

The findings from the evaluation of the first iteration led to some new design criteria to guide the design of the second iteration of the MusicJacket. These were:

- Feedback should have a user controlled on\off switch
- Vibrators should be placed on the body as well as the arms
- The vibrator-movement mappings that can be given at any one time should be kept to one
- Vibrators should be placed on bone rather than muscle.
- Where possible only one vibrator should be on each bone.

Changes

To meet these new criteria based upon the findings from the studies, the positions of some of the vibrators on the prototype were changed. The positions of the vibrators were intended to be easier to feel and distinguish and to suggest more intuitively what action the player should take to improve. This was intended to make the second iteration of the MusicJacket require less learning time than the original set up. First, the number of degrees of freedom which feedback was about was reduced to minimise the confusion caused by information overload from trying to distinguish and take action on vibrotactile feedback about several aspect of playing at once. Vibrotactile feedback would no longer be given about the height of the bowing hand since it is clear when a student has made a mistake in respect to this because they will get auditory feedback when they hit the next string. The vibrator that indicates when the violin is being held too high was also removed since it is more likely and more detrimental to playing for a student to let their violin drop too low than hold it too high.

Second, the other vibrators were moved to be positioned on the bones. The vibrator communicating the forward motion to the bowing arm was moved to the back of the right elbow, on the ‘funny bone’ as this was felt to give a more instinctive response. The vibrator which tells the user to lift up their violin was moved to the back of the left elbow near the ‘funny bone’ for the same reason. Third, two further vibrators were added to come on when the user needs to lift the violin to encourage them to lift their whole upper body rather than just their arm. These were placed on the sides of the torso on the lower ribs. The new setup is shown in Figure A-9.

Finally, users were given a way to switch off the vibrotactile feedback. Since players had their hands full with the violin a large pressure mat switch on the floor was used to give the user control over whether they receive feedback. This pressure mat was of the type used to sense intruders for burglar alarms. The pressure mat is linked to the power to the vibrators meaning that the user stands on the mat to switch the vibrotactile feedback on and when they want to switch it off they can simply step off the mat and the vibrotactile feedback will be switched off.

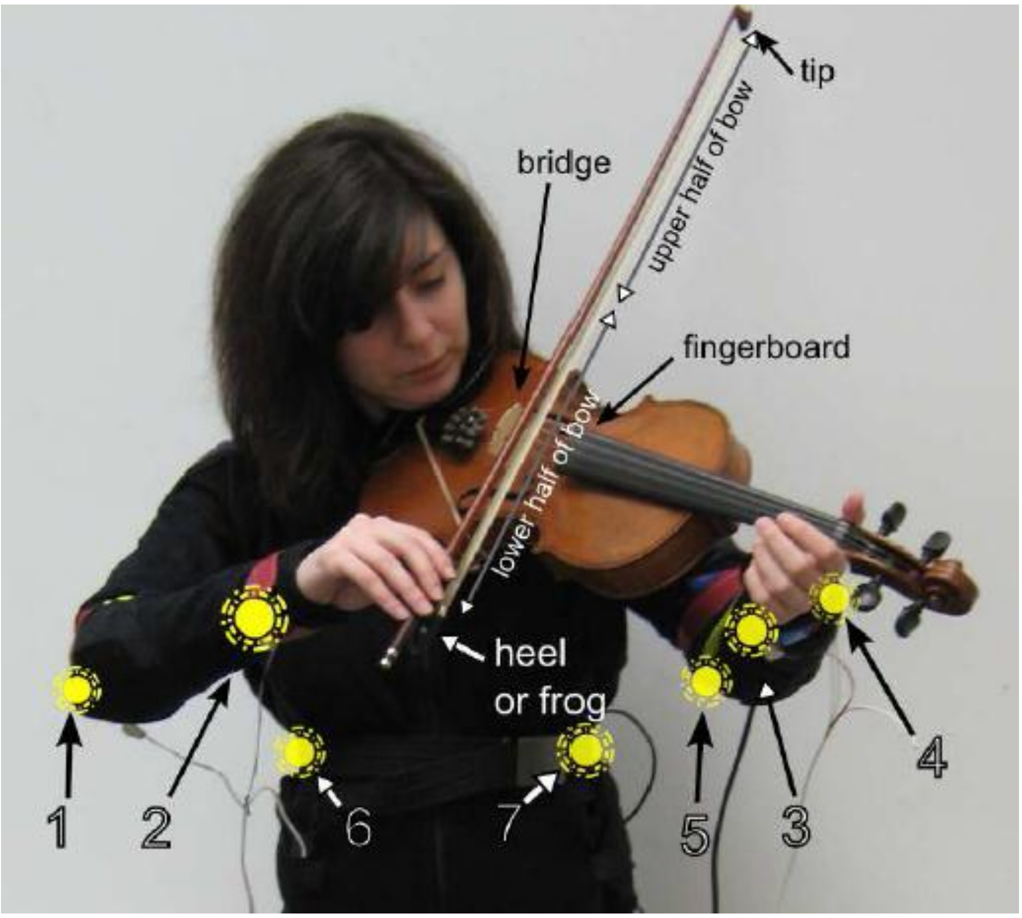

Figure A-9: Second prototype (1 move bow hand forward, 2 move bow hand back, 3 move violin left, 4 move violin right, 5, 6, 7 lift violin)

A.4 Discussion

The research described here was conducted primarily to gather information about how best to design the MusicJacket. However these findings also have implications for a more general understanding of real-time vibrotactile feedback. First, there are findings concerning information overload where participants find being given information about too many different aspect of playing overwhelming. Feedback which seemed reasonable at the design stage proved to be too complicated for use as real-time feedback all at once. This shows that the amount of information people can process in real-time while playing is limited. The limit appears to be reached when the vibrotactile feedback tries to give information about more than one aspect of playing.

Second, there is less spatial acuity in distinguishing between vibrators in real-time vibrotactile feedback than some of the literature about tactile instructions would suggest e.g. (McDaniel, Goldberg, Bala, Fakhri, & Panchanathan, 2012). Having feedback in real-time while playing the violin makes it harder for participants to accurately place where a vibration is coming from. We found that vibrators positioned on the same bone in arm were hard to tell apart while playing. The final design which is based on the user study and expert advice places the vibrators mainly on the joints. Interestingly these are also the anatomical reference points which Cholewiak and Collins (2003) found enabled better precision in identifying the location of vibrators. This makes the spacing of the vibrators for this application much wider than those used in some of the other vibrotactile systems (e.g. (Lieberman & Breazeal, 2007), (McDaniel, Goldberg, Bala, Fakhri, & Panchanathan, 2012)). It is also much wider than the 4cm resolution that is found in lab studies of sensitivity to pressure stimuli. This may be in part due to vibrations being conducted along the bones of the arm which would not happen in pressure stimuli. It may also be due to the distraction from playing the violin which means that the feedback is only being sensed peripherally.

Third, the exact position of the vibrator on the body is important. Subtly different positions on the anatomy can change the action that they suggest to people. A previous study which looked at intuitive responses to tactile stimuli (Spelmezan, Jacobs, Hilgers, & Borchers, 2009) only looked at large scale positioning of vibrators, for example comparing the front of the thigh to side of the thigh and focused on how different patterns of vibration on a particular limb can suggest particular movements e.g. vibrations running down the back of the thigh causes people to bend their legs. However we found that the fine-grained anatomical position of a single vibrator can affect what response is most intuitive. In particular, vibrators positioned on a muscle makes that muscle want to contract and was felt to be uncomfortable. Whereas positioning it on the same limb but on a bone gives a different kind of signal. This shows that even though spatial precision of perception may be lower than expected the positioning of vibrators is still very important. Rather than differentiation between vibrators in terms of their distance apart it is more important to think about their position with respect to the anatomy of the player.

Appendix B – Studying Real-time Vibrotactile Feedback Using the MusicJacket

B.1 Introduction

This appendix reports an investigation the efficacy of real-time vibrotactile feedback for learning and teaching the violin. The small scale study in appendix A suggested that vibrotactile feedback may support learning the violin, but this needs to be tested more rigorously now that important details such as the position of vibrators have been optimised. For this reason we conducted a controlled user study to compare groups of novices practicing with the feedback compared to those practicing without. This aimed to show whether real-time vibrotactile feedback can impact on the way learners progress and answer the following research question:

Can real-time vibrotactile feedback improve the way novices bow the violin?

B.2 Hypothesis

To address the research question a laboratory study was carried out with the MusicJacket. This used a traditional experimental methodology to test the following hypothesis:

Real-time vibrotactile feedback improves the violin bowing of adult beginner violin players.

The study was designed with a test group which learned with the feedback and a control group learned without the feedback. Motion capture data was collected from both of these groups. If the hypothesis is correct then significant differences should be found within the test group between data collected while using the feedback and without it. There should also be differences between the control group who had not used the feedback and the test group. This study was planned and analysed in collaboration with Erwin Schoonderwaldt and Jon Bird as well as my PhD supervisors.

B.3 Participants

Eight adult volunteers participated in the user study: five males and three females. The participants were all in their mid- to late-20s. Two participants had a strong musical background but did not have prior experience of violin playing; two participants had occasionally played the violin but had never received formal lessons; and the remaining four participants had no experience of playing any musical instrument. The participants were divided in two groups—a feedback group (coded as T1–4) and a control group (C1–4)—which were balanced with respect to the participants' musical experience and gender.

B.4 Method

Each participant met with the researcher(s) for six sessions. The first two sessions were devoted to learning the basics of playing the violin; namely: holding the violin, holding the bow, how to use the bow on an open string[3]. In the third session a measurement of their bowing movement was taken using the Animazoo motion capture suit (Synertial (formerly known as Animazoo), 2014), this is referred to as the *'before'* measurement. In the fourth and fifth sessions the participants in the test group practiced with real-time vibrotactile feedback from the MusicJacket, the participants in the control group practiced wearing the motion capture suit without feedback. During the fifth sessions measurement of these two groups were taken while playing without feedback, these are referred as *'between'*. An extra set of measurements was taken with the test group in this session which measured their playing while using the real-time vibrotactile feedback these are referred to as *'during'*. The sixth session took place at least one day after the last feedback session. In this no feedback was given to either group and an '*after'* measurement was taken using the motion capture suit. An overview of the sessions is given in Table B-1.

Each session ran using the same format. After putting on the motion capture jacket (and feedback if in the test group) the jacket would be calibrated by setting the reference stroke (as described in Appendix A.3.1). This was done with both control and test groups as the reference stroke is not only useful for giving feedback, it also acts to remind players what they are aiming for and is essential for analysing their bowing after measurements have been collected. Then the participant played a number of defined exercises. These are as follows:

1. Long notes using the full length of the bow from heel to tip
2. Short notes using half the bow from heel to the centre of the bow
3. Short notes using half the bow from centre of the bow to the tip
4. Mixed exercise using a variety of bow lengths in the form: long short short, long short short
5. *'Hot cross buns'* using a variety of bow lengths. The rhythm of this nursery rhyme was played on an open D string, ignoring the melody.

Each exercise was played for approximately a minute each, sometimes broken up into several shorter bursts. In the test group, participants would play each exercise initially without feedback and the feedback was switched on while they were playing after about 30 seconds. All sessions were videoed to capture observational data.

Session	Content	Measurement
1	Learning violin basics and exercises	N/A

[3] An open string is when you bow the string without any fingers down on the fingerboard

2	Learning violin basics and exercises	N/A
3	Practicing exercises (no feedback)	Before
4	Practicing exercises (with feedback for test group)	N/A
5	Practicing exercises (with feedback for test group)	Between and During
6	Practicing exercises (no feedback)	After

Table B-1: Overview of Sessions

B.5 Findings

To find out whether real-time vibrotactile feedback is effective in improving bowing we analysed four different aspects of the data. First, the measurements taken in the different sessions from the control and test groups were compared to examine how feedback influenced their progress and test the hypothesis put forward in the introduction to this study. Second, we analysed at the micro-level what happened to the bowing trajectory at the point that the feedback is switched on. Third, we looked at how straighter bowing was physically achieved and whether the feedback led to the correct arm movements. Fourth, we give a qualitative analysis of how the participants responded to the feedback in the study.

B.5.1 Within and Between Participants Comparison

The results from participants in the test group and the control group can be seen in Figure B-1. The data shows the deviation from the correct trajectory measured when participants are playing near the tip of the bow averaged over all the exercises that involved using the full length of bow (long bows, mixed exercise and '*hot cross buns'*). The deviation at the tip of the bow was chosen because this is the point where the deviation is likely to be largest because it is the most difficult part to keep straight.

The most striking pattern in this data is the contrast in the *before* measurements of the test group and the *during* measurements. This shows a significant reduction in deviation from the reference stroke ($p<0.01$) between their playing in session 3 before they had feedback and their playing with feedback in session 5. Moreover we can see a difference in the *between* measurement and the *during* measurement. This is interesting because both these measurements were obtained from the same session and show clearly the difference between playing with feedback and playing without. This supports the hypothesis that real-time vibrotactile feedback can improve the bowing of adult beginners.

There are no significant differences between the data from the control group and the test group. This is perhaps to be expected given the small sample size. However trends can be observed in the data which suggest that the test group have progressed differently compared to the control group. The data shows that all participants in the feedback group improved their bowing by playing closer to the reference trajectory in the *between* measurement (playing without feedback but having practiced with the feedback) compared to the *before* measurement; whereas, in the control group, only two participants moved toward the reference stroke. For one of the participants in the control group, C3, this change does not appear to be the product of gradual learning so much as a complete shift in the bowing style. The *after* measurements show that the improvements of the test group were maintained some days after by two of the participants (T1, T3). For participant T2, there was no clear difference between the *after* and *before* measurements, confirming the impression of the experimenters that he was falling back into his old ways. Participant T1 showed the largest deviation from the reference bowing path during the last session, which might, however, be partly explained by a slightly skewed bowing reference line. In the control group C1 and C3 both show some improvement compare with their initial playing, C1 in particular showing gradual improvement over the study, whereas C2 and C4 show no improvement.

In sum this data shows that during the time playing with feedback participants played closer to the reference stroke, suggesting that real-time vibrotactile feedback may help people to progress faster than those practicing bowing without feedback.

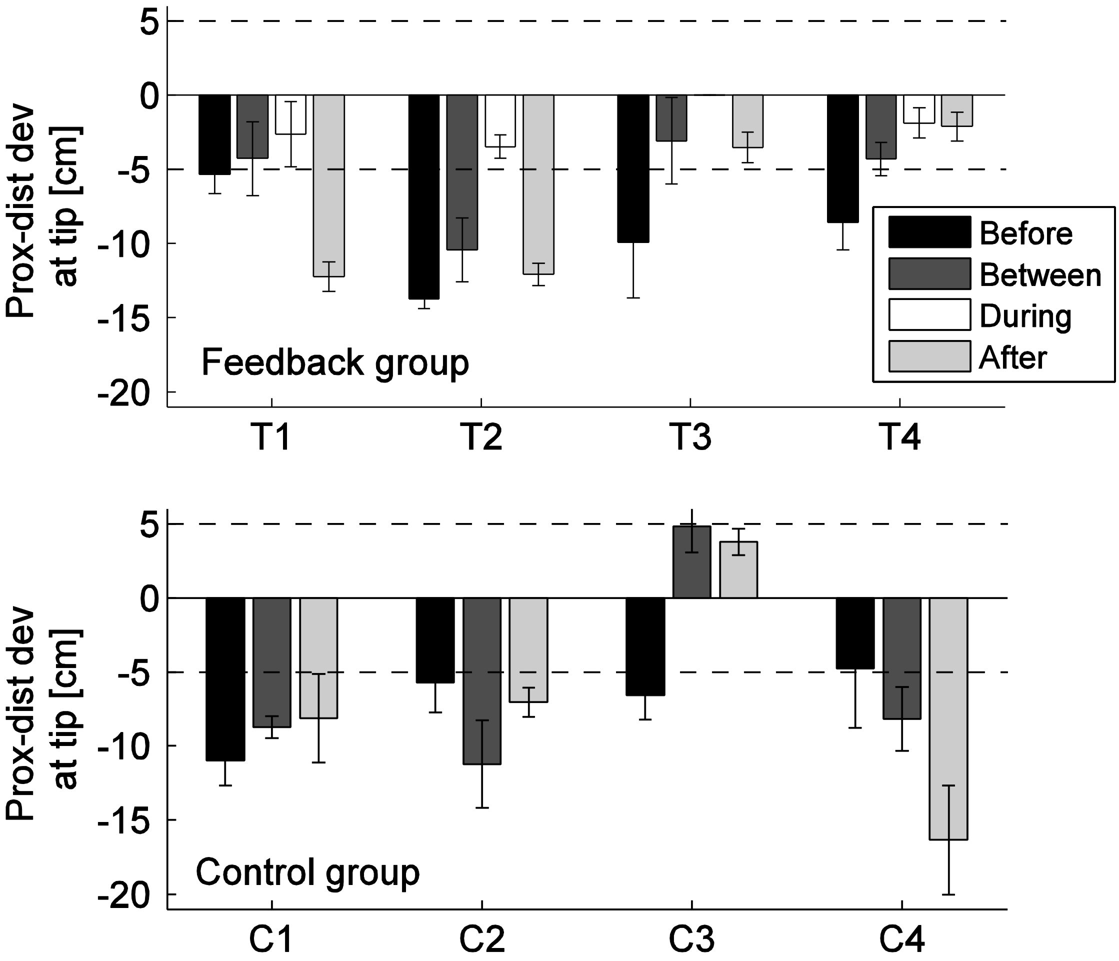

Figure B-1: Average deviation from the ideal bow stroke measured at the tip of the bow. T3 has no during measurement as this recording was not able to be made at the time of her participation. The error bars represent 95% confidence intervals from the mean.

B.5.2 Micro-level Analysis of the Guiding Effect of Vibrotactile Feedback

Next we examine in detail each test participants' bowing trajectory before the feedback and how this was adjusted in the moments after feedback was switched on (see Figure B-2, Figure B-3, Figure B-4 and Figure B-5). We do not have this data for participant T3 as the MusicJacket did not record when it was switched on for her but we do examine other data from her in a later section. Participants T2 and T4 reacted to the feedback by interrupting their established bowing pattern and realigning their bow. T1 was more gradual in her response to the feedback and gradually adjusted her movement over many bow strokes. More detail is given about these two types of response in the following two sub-sections.

Interrupted bow stroke

Figure B-2 shows T2 playing an exercise which required him to play a pattern of long and short bow strokes. T2 began the exercise without feedback (dashed lines) and after 38 seconds of playing the researcher switched the feedback on (solid line). During the initial playing the participant established a natural bow stroke which was not very close to the ideal stroke (deviating by up to 20cm from the reference stroke). When the feedback was switched on the deviation from the reference stroke triggered the vibrotactile feedback (thick line). Figure B-2 shows that the participant responded to this by interrupting their natural stroke and moving their trajectory towards the reference line. We see approximately 10 seconds of confusion and uneven bowing before the participant is able to return to the rhythm of the exercise (long short short). This method was very effective in reducing the amount of feedback he triggered, as he only triggered a lot of feedback on the first two bow strokes and then barely anything after that.

When the feedback was switched on for T4 he also interrupted his bowing rhythm to adjust his position and then returns to rhythmic playing once he has found the bowing trajectory that does not trigger feedback (Figure B-3). For him there is about 30 seconds of confusion before he finds the right place. This is an example from his first session with the feedback. By his third session with the feedback he no longer needed to interrupt his stroke when the feedback was switched on and was able to make the minor adjustments suggested by the feedback while keeping up his rhythm (see Figure B-4). This is partly because he had learnt how to respond to the feedback quicker and partly because as he progressed through the sessions his starting bow stroke moved closer to the reference trajectory so that he only needed to make small adjustments.

Gradual adjustment

Participant T1 reacted to the feedback in a different way. She did not interrupt her playing; instead she gradually adjusted her playing avoid the feedback over many bow strokes. Figure B-5 shows that it takes 30 seconds before she succeeds in reducing the amount feedback she is triggering and even then she is still triggering some. This approach of gradual adjustment is one which she consistently uses through all her sessions with feedback. Although it does not improve the bowing trajectory as quickly or as much as the approach of interrupting playing and realigning, the aggregated data shows that it is still effective for improving bowing and has the advantage that it would not interfere so much with the flow of a music practice.

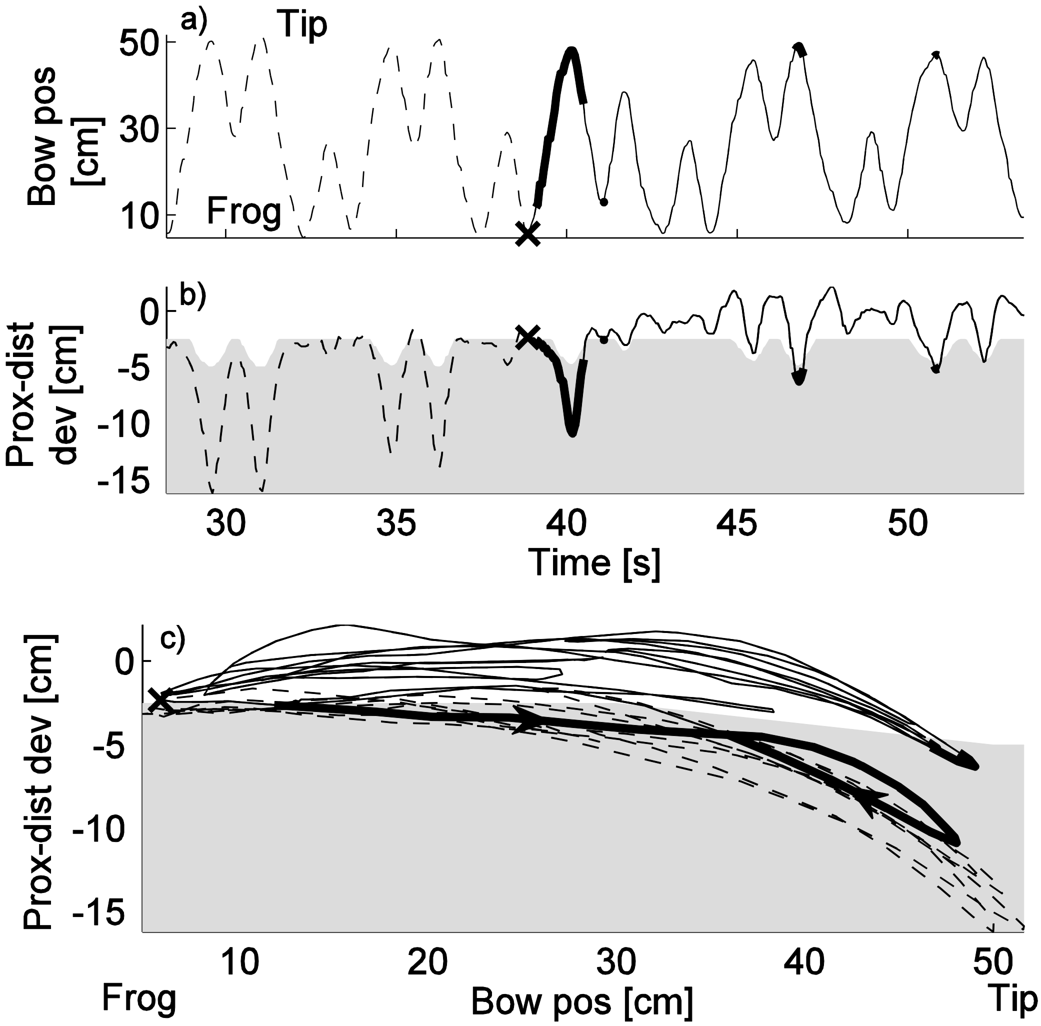

Figure B-2: The change in T2's bowing movement as feedback is introduced. The dotted lines show the participant playing with the MusicJacket switched off. The X shows the point when the MusicJacket was switch on and the solid lines shows the movement when the MusicJacket is on. The lines that are bold show when vibrators are actually active. The shaded regions indicate the zones where vibrotactile feedback would be triggered if the MusicJacket was switched on.

a) Shows the movement of the bowing hand perpendicular to the strings over time.

b) Shows the deviation from the reference bow stroke over time (in the direction parallel to the violin strings). Ideally this would continuously be zero.

c) Shows the movement of the bowing hand as viewed from above. Imagine that the strings of the violin are running parallel to the y-axis. The reference stroke is y = 0.

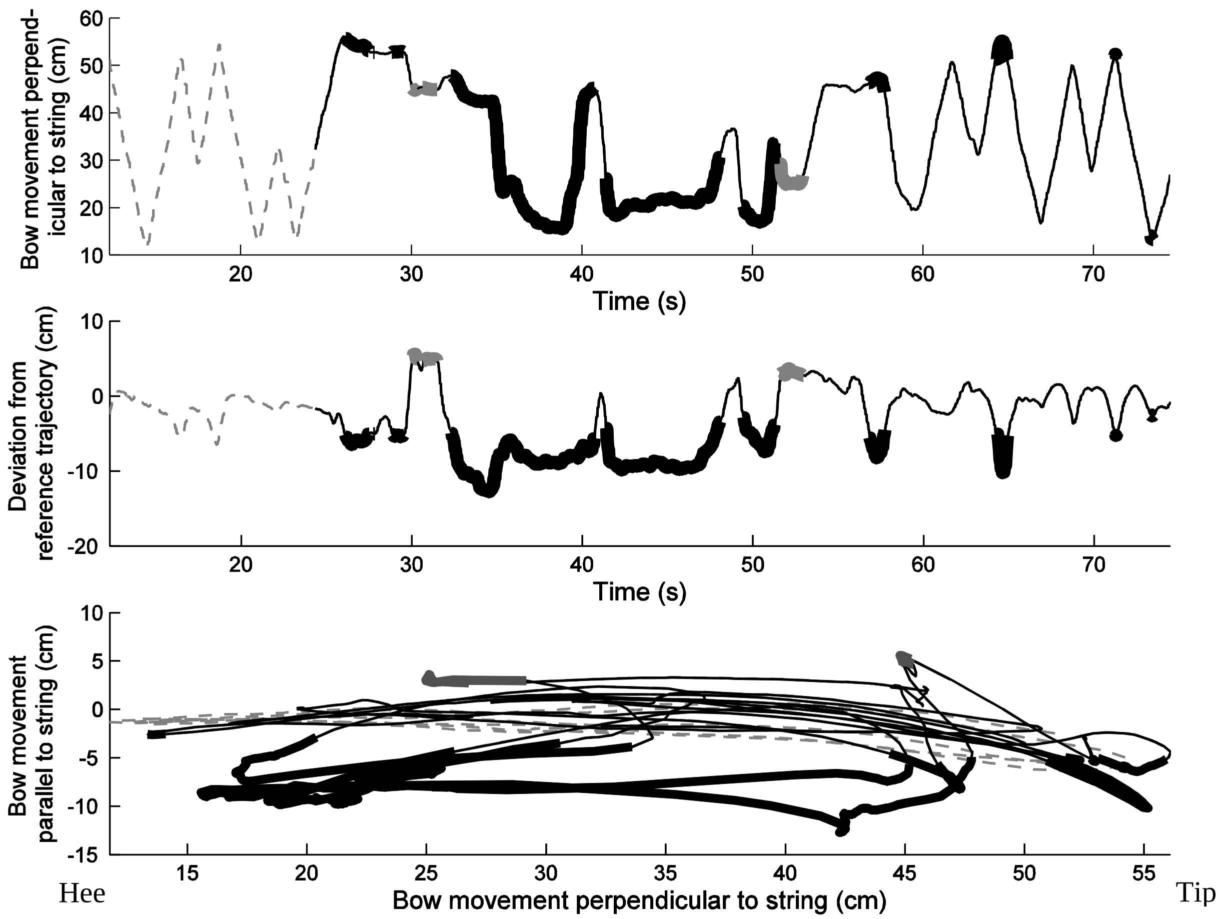

Figure B-3: The change in T4's bowing when feedback was switched on in session 4 (first session with feedback). The dotted lines show the participant playing with the MusicJacket switched off. The solid lines show the movement when the MusicJacket is on. The lines that are bold show when vibrators are actually active: black is the vibrator on the elbow, grey is the vibrator on the wrist.

a) Shows the movement of the bowing hand perpendicular to the strings over time.

b) Shows the deviation from the reference bow stroke over time (in the direction parallel to the violin strings). Ideally this would continuously be zero.

c) Shows the movement of the bowing hand as viewed from above. Imagine that the strings of the violin are running parallel to the y-axis. The reference stroke is y = 0.

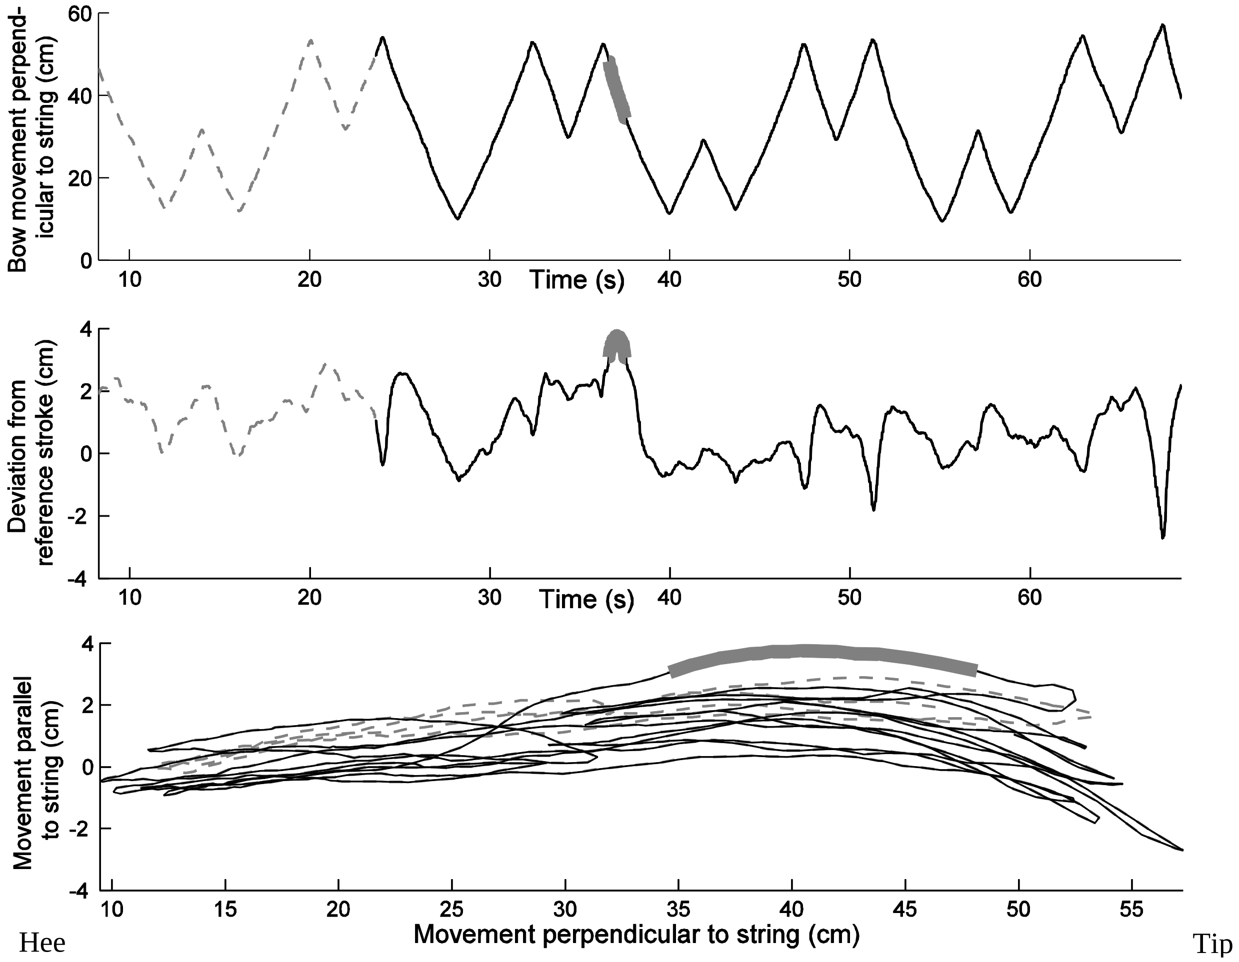

Figure B-4: The change in T4's bowing when feedback was switched on in session 6 (last session with feedback). The dotted lines show the participant playing with the MusicJacket switched off. The solid lines show the movement when the MusicJacket is on. The lines that are bold show when vibrators are actually active: grey is the vibrator on the wrist.

a) Shows the movement of the bowing hand perpendicular to the strings over time.

b) Shows the deviation from the reference bow stroke over time (in the direction parallel to the violin strings). Ideally this would continuously be zero.

c) Shows the movement of the bowing hand as viewed from above. Imagine that the strings of the violin are running parallel to the y-axis. The reference stroke is y = 0.

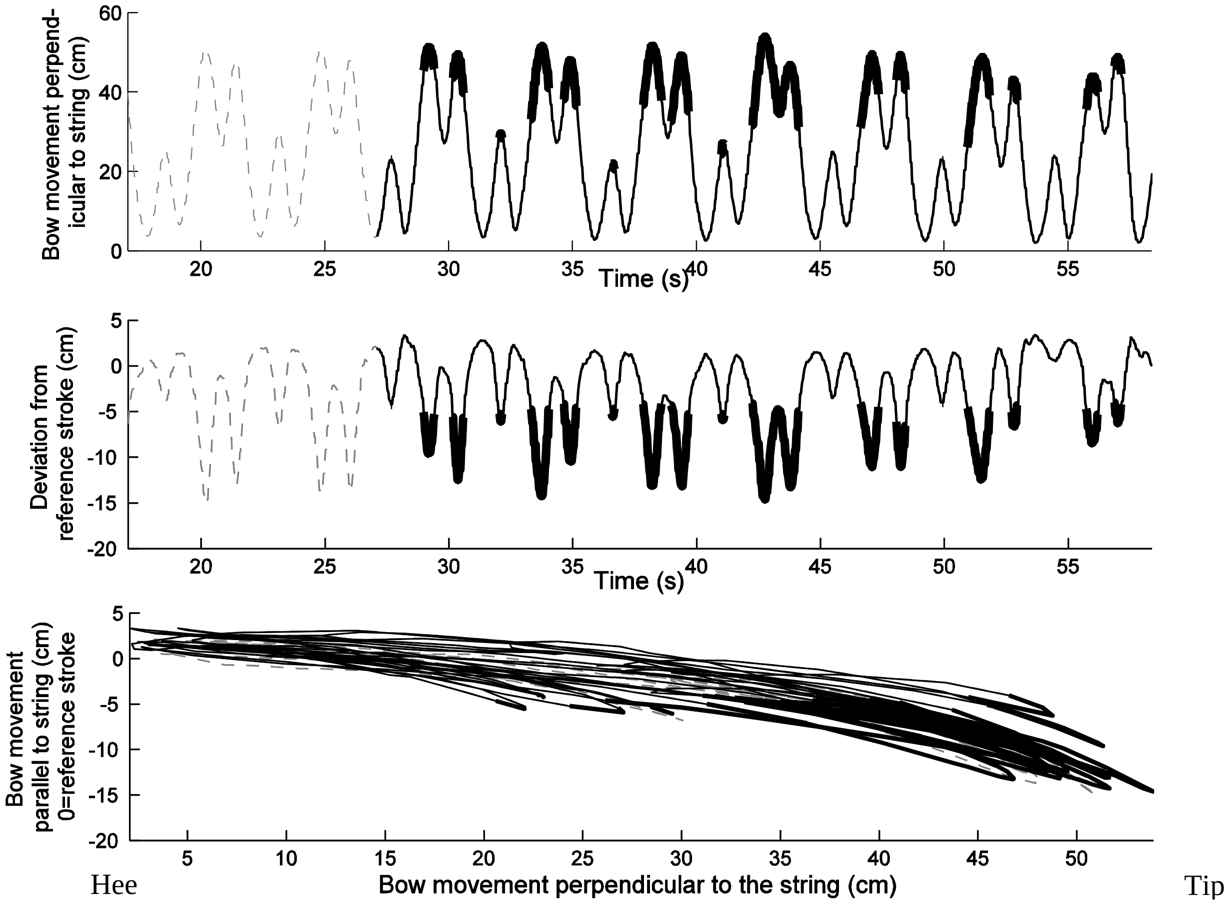

Figure B-5: The change in T1's bowing when feedback was switched on. The dotted lines show the participant playing with the MusicJacket switched off. The solid lines show the movement when the MusicJacket is on. The lines that are bold show when vibrators are actually active: black is the vibrator on the elbow.

a) Shows the movement of the bowing hand perpendicular to the strings over time.
b) Shows the deviation from the reference bow stroke over time (in the direction parallel to the violin strings). Ideally this would continuously be zero.
c) Shows the movement of the bowing hand as viewed from above. Imagine that the strings of the violin are running parallel to the y-axis. The reference stroke is y = 0.

B.5.3 Changing Movement

To play a straight bow involves a small movement from the shoulder joint and a larger movement opening up from the elbow (Figure B-6). Often novices fail to make this movement from the elbow and only move from the shoulder resulting in a curved bowing trajectory. Figure B-7 shows the difference in terms of joint angles between an expert bowing and a novice bowing as demonstrated by an experienced teacher. In this section we look at the findings from a special participant (T3) who observations during in the sessions showed had the particular problem of only moving her arm using her shoulder joint and not her elbow (see Figure B-8a).

Consequently she made the most pronounced progress of all the participants from the *before* measurement to the *between.*

Examining the arm angle data of T3 we can see that she has changed the way she moves her bowing arm after having used the vibrotactile feedback. Figure B-8a shows that in the *before* session she is moving her arm mainly from the shoulder, like the teacher's demonstration of novice bowing. In the *between* session (after using the feedback, Figure B-8b) her bowing has transformed and she is now largely moving from the elbow, just like the teacher's demonstration of expert bowing. This shows that in her case feedback didn't just make her play closer to the reference stroke; it also made her move her arm in the recommended way. This demonstrates that for players that have an individual difficulty with a certain movement real-time feedback can be especially helpful.

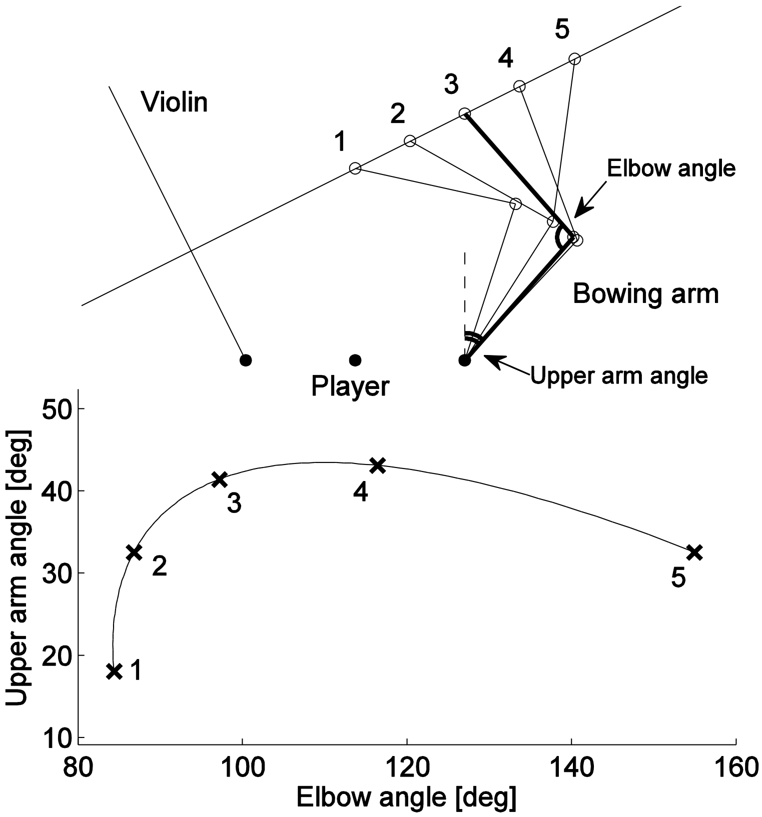

Figure B-6: Model of the required movement of the elbow and shoulder to play a straight bow. The player moves their arm from the shoulder first and then opens out their elbow to keep the bow moving in a straight line perpendicular to the strings.

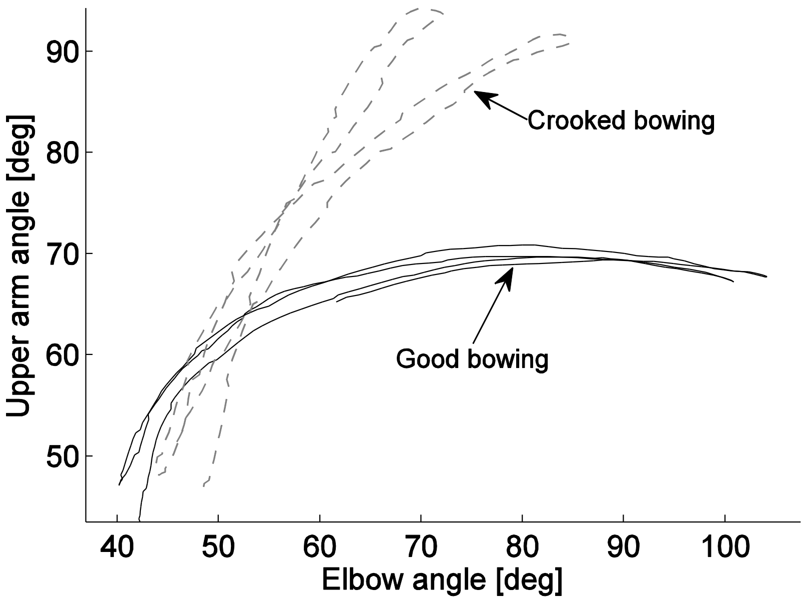

Figure B-7: Angle of the joints for crooked and expert bowing, as demonstrated by a violin teacher. Expert bowing involves more use of the elbow joint than novice bowing which relies more on the shoulder (larger changes in the upper arm angle).

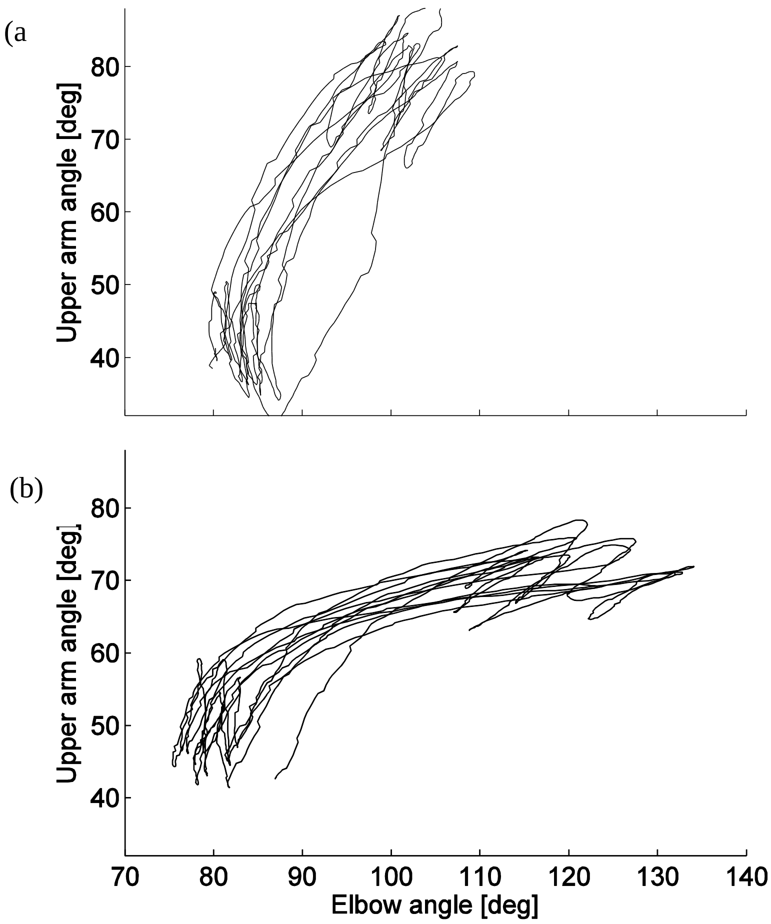

Figure B-8: T3's joint angles *before* feedback – top; and after some feedback (*between*) – bottom. Compare these to the Figure 6 and notice the clear progression from the novice to expert bowing movement.

B.5.4 User Response to Feedback

In the sessions with the vibrotactile feedback some of the participants talked about the movement of their bowing arm, particularly about changing the angle of the hand, fingers, or wrist and a need to extend the elbow. This aspect of body awareness was not expressed in the same way by any participants in the control group. Most test participants also noted that they experienced most feedback at the extremities of the bow—either the heel or the tip. This meant that they knew where they had to focus their attention when practicing. An unforeseen and undesirable consequence of this was that some participants chose to reduce the amount of bow they played to avoid feedback, rather than changing the angle of their bow at the tip or heel. One of the control participants also mentioned that it was harder to play near the tip, showing that although the feedback makes problem areas more obvious, it is still possible to notice them without its help. In general, the participants seemed to think that the feedback was helpful, and one, in particular, said that *"if I was to keep playing, I would choose with feedback ... because it's useful."*

B.6 Extension to study

After the experiment was completed, we investigated further to see how the participants in the control group would react to the feedback after having had more time to establish their own "natural" behaviour. Therefore the participants in the control group had an extra session where they used the vibrotactile feedback. The same procedure was used as for the participants in the test group. During the initial part of each exercise, there was no feedback, and feedback was switched on after the participants had developed a stable bowing pattern. The results in Figure B-9 show that three of the four control participants significantly reduced their deviation from the reference trajectory with the feedback during this session. The participants were able to adapt to the feedback quickly, as this improvement happened in a single session; whereas the participants in the main study had been given a whole session to practice with the feedback before measurements were taken. One participant, C3, did not improve, this is because he was already playing so close to the reference stroke that the feedback would not have been activated. The data from these control participants provides further evidence for the efficacy of real-time vibrotactile feedback to change people's established bowing pattern and move it towards the reference stroke. Moreover it demonstrates that for novices who have been focusing on bowing, the positions of the vibration motors must be quick to learn how to use to give this rate of improvement in a single session.

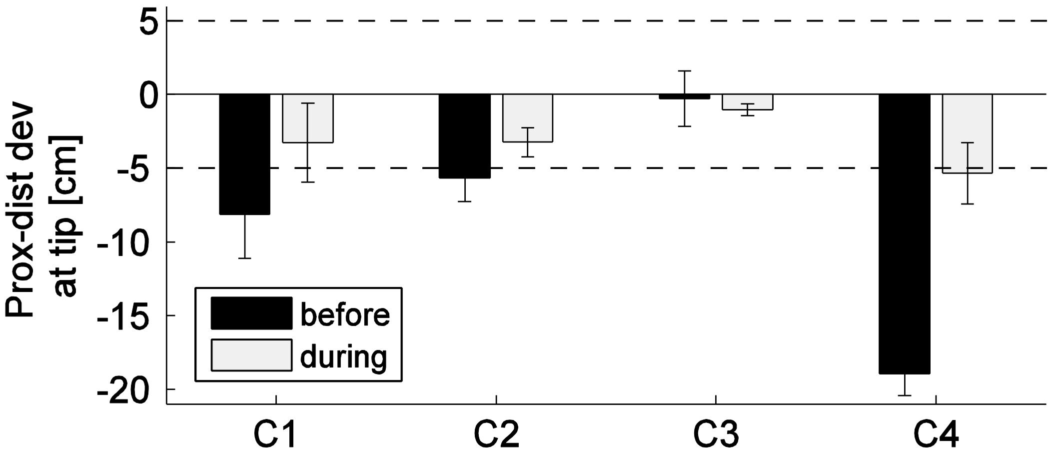

Figure B-9: Extra measurements from the control group using the feedback. The bars show the average deviation from the reference bow when playing near to the tip of the bow (the hardest place to stay on track).

B.7 Discussion

These finding show that real-time vibrotactile feedback is effective for improving bowing in a way that can be quantitatively measured. It is interesting now to consider how the vibrations are acting on the player to change their movement. The user verbal responses to the feedback show participants being more aware of their body movement and the angle of their joints. This suggests that the vibrations on the body maybe acting to enhance awareness of proprioception. From the stand point of embodied interaction (Dourish, 2001) this would be an advantage since the body plays a key role in the way we understand the world. Studies with dancers (Kirsh, 2010) show that the process of marking which is a combination of mental practice and smaller movements is superior for learning a technically demanding piece than mental practice alone or the full movements . This shows that to maximise learning, mind and body should work together. This is why these verbal responses showing participants thinking more about their body are promising. It suggests that the vibrotactile feedback on the body is drawing attention to the role of the body in playing the violin and encourages players to think about their body and understand their playing through their body.

The microlevel analysis shows two ways that participants can react to the feedback. One is to interrupt their playing and then realign their bow to a better trajectory. This can be thought of as the feedback attracting the participant's attention to a problem and then once they are focused on the problem, helping them to improve. In this scenario the player is making the adjustments with the feedback as their central focus. The other reaction is for participants to keep playing and very gradually change their bowing. In this scenario the feedback is still attracting attention to a problem but it is likely that this attention is more peripheral. The case of T4 shows that a participant can change in their response to feedback as they get more familiar with the feedback and they only need to make minor improvements to their trajectory. The aim of the real-time vibrotactile feedback at the design stage was that it would be peripheral and allow players to continue with their normal practice without interruption. This analysis has shown that it can play this role in some circumstances, particularly when participants are familiar with it and are at a level where they only need minor improvement. In addition, in some cases where participants need to make large changes to their bowing trajectory, it can attract a central focus and help the player make a large improvement. Given the extent of the improvement in these cases this use of the central focus may be well justified and once the bow has been realigned, players can return to focusing on other elements of playing.

The feedback from the MusicJacket can be thought of as negative feedback because it points out mistakes rather than affirming achievements. This raises questions about whether will be motiving for students, or whether participants will feel negative about having their failings pointed out in this way. The participants in this study did not make any comments about this.

However they were only learning to play the violin for the purpose of this study and therefore their motivation to improve was not based on long-term goal to become a violinist. The effect of negative feedback may different when it comes to real learners who have chosen to learn the violin. Having shown in this study that real-time vibrotactile feedback can be effective for helping complete beginners to learn the violin, we were interested to start exploring whether the MusicJacket can be useful to genuine learners who are already attending lessons (this is explored in more detail in Appendix C).

Appendix C – Using MusicJacket in the Classroom

C.1 Introduction

Having shown in the previous study (appendix B) that the MusicJacket can improve the playing of complete beginners in a controlled experiment, we were keen to investigate whether this success could be translated to real learners in their day to day lessons. The aim of this study was to examine the efficacy of the MusicJacket as a form of feedback in-the-wild, for children who are already learning the violin. By doing this it aims to investigate the following research questions:

Does real-time vibrotactile feedback aid the learning and teaching of bowing and posture for children already learning the violin?

How can real-time vibrotactile feedback be incorporated into violin lessons so that it facilitates learning?

This study differs from the earlier one, because it focuses on the practical usefulness of real-time vibrotactile feedback for teachers and children in the classroom. To do this we examine how the MusicJacket was appropriated in a real-life context as this finds out how people really choose to use a new device not just how they think they'll use it or should be using it (Rogers, 2011). This is particularly important to investigate since pupils and teachers will have had very little experience with any kind of real-time feedback so cannot be expected to understand how they might use a device like the MusicJacket if it is given to them out of context. It is important to study the MusicJacket with real violin students in their normal activities from a motivational point of view as well. Students already learning the violin have shown that they are interested and motivated in learning and will be more driven than beginner participants to use the MusicJacket to improve their playing because they care about the outcome for their playing. Because they are likely to intrinsically motivated to play the violin irrespective of the study, we will also be able to see how the real-time feedback interacts with this, whether it adds and exciting challenge to the their playing or make it seem more arduous.

There are also strong arguments to study the use of the MusicJacket over the long term. There is evidence that after a certain period of time the novelty aspect of a new technology wears off and that after longer use people appreciate different aspects of a technology than they do on first encountering it (Karapanos, Zimmerman, & Martens, 2009). The MusicJacket may initially have a strong 'wow' factor, especially for children, but we want to move beyond that, and explore hoe the technology works over an extended period of time in a real world teaching setting. Hoggan and Brewster (2010) also stress the importance of longitudinal studies in situ for mobile devices with multimodal interfaces, since certain modalities may turn out to work better in some environments and, of particular pertinence to our study, people may develop different habits under different workloads. These forms of usage cannot be uncovered in short duration laboratory experiments. Furthermore, a long term study in the classroom allows all parties (teachers, pupils and researchers) to get to know each other, as well as the technology, and thus build up a relation of trust, from which a process of proper exploration can flourish. This means that both teachers and pupils can become actively involved in exploring how best to use a technology – as opposed to the more passive role they would have in a controlled experiment. Our approach was to let teachers and pupils appropriate the MusicJacket and over the course of a number of lessons make use of the real-time vibrotactile feedback in ways best suited to their individual needs.

C.2 Participants

Two teachers (A and B) took part in the study, both of whom had previously been involved in informing earlier stages of the design process (see Section A.2). Both teachers had many years of experience teaching the violin. The teachers selected pupils to participate in the study who they thought might be most interested in using the MusicJacket or might benefit most from what the MusicJacket had to offer. One teacher chose a group of six children (group A) aged between six and nine which she regularly taught at a local primary school. The other teacher chose four older students aged between 10 and 13 (group B) who she taught in private lessons.

Working with children, particularly the younger ones, led to a very different set of demands from the MusicJacket and our approach to studying it compared to the adults in the earlier studies. First a smaller jacket needed to be made as the original MusicJacket was much too large. The children in group A also got tired very quickly; they could not hold the violin up for more than 10 minutes. It is also hard for young children to explain what they are feeling while using the MusicJacket, so we resorted to close examination of their facial expressions.

C.2.1 Group A - The Primary School Children

This group consisted of six girls who had started their violin lessons only a few months prior to the start of the study. With the exception of one girl aged nine, these children were just six years old. Their violin lessons took place during school hours in a spacious music room at their primary school. Their teacher was employed by the local Music Service, and visited a number of schools in the area to teach violin, as well as giving private lessons to children and adults. She had selected these children as they were absolute beginners and because she felt the children, their parents and the school would be keen to participate.

The children were normally taught as a group for a twenty minute session. When we observed such a lesson they were in the beginning stages of learning how to hold their violin. The lesson involved singing and dancing to help children get to grips with the basics. The children sang and danced to music at the start of the lesson to loosen their muscles. Then they did actions to a song about an elephant to help them put their violin on their shoulder: this involved holding their violin up to their nose like a trunk before bringing it down onto their shoulder. The teacher used a story about a mouse to help the children practice holding the bow. A CD was played and the children plucked their strings in time with the tunes. The children only spent a short amount of time playing the violin (approximately 10 minutes in total) because they found it so tiring. The teacher provided a playful environment in which active movement was very important and children showed that they were enjoying themselves by smiling and giggling as they joined in. By the time our study started they had just moved on to bowing, but this was still in the early stages. These observations are important because it shows how the children were motivated by the fun environment that they were taught in and part of the test of the MusicJacket would be how well it fitted into this environment.

C.2.2 Group B - The Older Children

A group of older children also took part. The four children in this group were 10 to 13 years of age and had all been playing for about three years. Pupils were used to having individual lessons – one of them had private lessons outside school hours, but for the others the teaching was organized during school time. Their teacher also had many years of experience in violin teaching. The teacher had selected these pupils as each of them had trouble with a particular aspect of bowing and she felt they could all benefit from focused bowing sessions.

These lessons were less playful and more directly focused on violin playing than the lessons of the six year-olds in group A. These children also had more stamina and were able to play for longer because they had been regularly playing for several years. However, with these older children, particularly the ten year-old, it was still challenging to encourage them to express their feelings about the MusicJacket and for this reason we also studied closely their facial expressions and body language to try and understand how they were feeling.

C.3 Method

The study with the MusicJacket ran over five session spread over two months. In each session each participant had an individual lesson with their teacher with two or three researchers including me also present. At least two researchers were necessary to efficiently operate the MusicJacket, one researcher helping the child to put it on and another controlling it from the laptop. Having multiple researchers present at these sessions was also useful for taking observations and checking them against one another. These lessons were in addition to their regular lessons and had the particular focus of improving bowing (group B) or violin hold (group A). We met with the primary school children during school hours in a spacious music room where their lessons were normally held. The sessions with the older students were held in the living room of one of the researchers. In the first session, all the children played wearing the motion capture suit without any feedback, this was to take a baseline measurement. In the subsequent sessions the feedback about violin hold and bowing was gradually introduced. A gradual approach was taken to make sure that the children had time to learn how to use the feedback and so that we could learn what was an appropriate level of feedback for the children so that they felt confident and not overwhelmed. Group A mainly focused on violin hold and group B focused on bowing, this decision was mainly made by their teachers who chose the feedback that was most relevant to their current stage of learning. Each session was videoed and written observations were also taken at the time. The motion capture suit was also used to record the player’s movements in every session. Each session had three phases: preparation, exploring the technology and playing exercises and pieces.

(i) Preparation

At the start the sessions in a particular location the equipment was set up, which took about 40 minutes. The pupil was then dressed in the MusicJacket, taking about 10 minutes, followed by the actual lesson which ranged from 10 minutes to half an hour depending on the age of the participant (group B had longer lessons that group A). It took only 1 minute to remove the system at the end of the session.

(ii) Exploring the Technology

Each lesson had a component in which children were encouraged to physically explore the technology. During the first session we introduced the concept of motion capture and encouraged children to look at themselves using the Animazoo viewer software (Synertial (formerly known as Animazoo), 2014) which shows a stick figure which moves in real-time with the wearer of the suit. Children were encouraged to move their arms, rotate their wrists, pretend to bow and make big arm movements. This exercise aimed to encourage the children to become more aware of their body movements. This exercise was formulated with the help of the Alexander Technique Specialist who we consulted in A.3.2 who thought that it would be particularly useful for very young children who tend to have very little understanding of how their bodies work.

In subsequent weeks we introduced different levels of vibrotactile feedback by switching on each actuator in turn. Each child experienced how by moving their bowing arm in the right direction they could switch off the vibrotactile feedback and that this indicates that they are now in the correct position. This exercise was repeated most weeks.

(iii) Exercises and Pieces

The actual exercises that were carried out changed during the course of the study. Initially, we started with the same exercises we had used in Study 1. However, the exercises that worked well with adult novice players in a lab setting were not appropriate for the participants in this study. They were too difficult for the young children in group A who had only just started using the bow and were used to a far more playful style of teaching. They were too easy for the older children in group B who were used to playing more sophisticated pieces. In consultation with the teacher we quickly adapted our approach and got the children to play the pieces they were studying at that moment.

By week 3 we also established that the concept of straight bowing was not appropriate for the primary school children (group A), as their current bowing was not fully developed. However, the issue of 'holding up the violin' was very relevant for this group. This was the reverse of the other group whose priority was to improve their straight bowing.

We gradually developed a programme in which the early part of the lesson contained an element of 'deliberately playing badly', in which the pupil was encouraged to deliberately play in the wrong way in order to feel the feedback that resulted. This was done using an elementary exercise – bowing a few strokes for the younger children, or playing a scale for the older children. This is an example of an exercise in which the feedback is in the foreground and pupils are encouraged to fully concentrate on it. Subsequently we explored tasks where the focus was not the feedback itself. For example, in group B the teacher would give normal instructions on some aspect of their playing (say rhythm, or intonation) while in the background the MusicJacket was providing real-time feedback on the children's bowing.

C.3.1 Data Collection

In each session, the motion capture data was recorded while the children played with or without feedback. The researchers observed what happened in each lesson and made notes. Researchers also asked pupils questions at different points about whether they had felt the feedback, whether they knew how to respond to it and what it felt like. Lessons were videoed and after the study many were transcribed. In cases where a child did not express themselves very much verbally close attention was paid to their facial expressions and other non-verbal forms of communication – in particular comparing the facial expressions of the same child in different circumstances, such with or without feedback. This approach was necessary due to the age of the children as it can be difficult for young children to express themselves clearly to an adult especially when the adult is a researcher who may be perceived as having a vested interest in showing that the MusicJacket is perfect. After the study both teachers were interviewed about the study.

C.3.2 Releasing Control

Over the course of the study our methodology evolved to meet the challenges and opportunities that arose. Initially this study had been planned to produce a similar type of data to study 1. Namely, we split the children into a control group and a test group and only the test group was given feedback. We planned to use the same exercises as study 1 so that all participants could be compared playing the same exercises. However we found that this methodology was inappropriate for an in-the-wild study. The exercises were not engaging or the correct level of challenge for the participants. The young children in group A had only just begun learning to use the bow and the difficulty of the long bowing exercises meant that the teacher had to guide their movements in order for them to be able to carry out the exercises. This guidance would negate the effect of the feedback therefore the exercises needed to be changed. We took advice from their teacher and for the rest of the sessions they played the pieces they were currently working on. With the advanced children in group B the reverse was true. The exercises were quite simple for them and could be performed quite mechanically without triggering much feedback. However when they started playing their normal pieces from their regular lessons we observed that they had more difficulty keeping their bow straight and had more need of the feedback. Therefore with this group we kept the planned exercises but also asked them to play some of their normal repertoire as well.

In response to these observations we adapted our method so that each student played their pieces they were currently working on in their normal lessons. This meant that we were studying them playing something that they had more motivation to improve and which was challenging for them not only in terms of bowing but also the other skills involved in playing violin. We also gave the teachers freedom to start introducing exercises that they felt were relevant. For example, teacher B asked her pupils to practice straight bowing by “bunny hopping” the bow up and down the string and playing with their eyes shut. By doing this we were able to see the feedback being appropriated by the teachers and incorporated into the teaching strategies they would normally use to teach bowing and violin hold.

As the study progressed it became clear that there was no way we could objectively compare the control group with the test group even though we had tried to balance them with respect to ability. The playing styles and abilities were so diverse and complex. Some of the participants which had initially been labelled as ‘good’ turned out to have quite rounded bowing but had other strong points such as good bow speed or bow hold which resulted in a pleasant sound. Whereas others, who had been initially thought of as ‘bad’, bowed straighter because they bowed slowly and had more time to focus on their trajectory but still made a less smooth sound because of the low bow speed. When looking at straight bowing, bow speed is important since it is not fair to compare the straightness of one bowing trajectory to another if they were occurring at very different speeds. Moreover each child had a different set of strengths and weaknesses and the children were playing different pieces, sometimes from memory and sometimes from written music, all of which posed diverse challenges. Therefore after the first three sessions the children from the control group also started to use the real-time vibrotactile feedback from MusicJacket. This gave a wider variety of individual cases to examine in terms of how they appropriated the real-time vibrotactile feedback in their playing. Unfortunately one of the participants in group B was unable to attend these later sessions so no data was gathered with her using the feedback straight bowing.

C.3.3 Adapting the MusicJacket to Individuals

During the study we observed that one of the students in group B had a different problem with bowing, namely, she did not use enough bow. This was a concern for her teacher and she dedicated a large amount of the lesson time to encouraging her pupil to use more bow. In a brainstorming session with the teacher and the pupil during the fourth lesson, we explored the possibility of adapting the MusicJacket system to suit the particular needs of this pupil. The following week we were able to present an adapted version which works by providing vibrotactile feedback at each end of the bow length if the student moves far enough along the bow. She was told to *aim* to feel the vibration as opposed to *avoiding* vibration in the usual set-up. By making this alteration we were able to make the MusicJacket more relevant to the needs of this pupil. This fitted with our objective of exploring how vibrotactile feedback can be usefully applied to real teaching situations. We did not make any other adjustments to the MusicJacket to meet other individual needs because we did not observe any other cases where adjustments were necessary to make the feedback relevant to a participant’s learning goals.

C.4 Findings

As already discussed in the method section, each participant was very individual, particularly the older children who had had more time to progress and develop their own individual style as well as bad habits. Therefore we have approached the analysis of the data from this study on a case by case basis, grouping together students when they showed similarities in the way they learnt with the feedback and separating out special cases. We draw on observations, video data, teacher and student comments and motion capture graphs to describe qualitatively the four main areas of findings. These are:

i. **Holding up the violin**, which was the focus for the children from group A
ii. **Straight bowing,** which was the focus for two of the children in group B
iii. **Using more bow,** which was the focus for one of the children in group B
iv. **Calibrating the MusicJacket** to personalise the feedback, which teacher B became actively involved in

i. Holding up the violin - Group A

The feedback from the teacher and the observations during the sessions showed that feedback about holding up the violin was the most successful type of feedback for group A, the primary school children. These children were still at stage where they needed to practice bringing the violin up to their neck and positioning it in the correct place. It was not unusual to see the violin slip back out from underneath a child's chin whilst they were playing, so any teaching approach which encourages holding the violin up was particularly relevant to them at this stage of learning.

The children learnt to respond to the vibration on their elbow and sides by lifting up their violin generally within the space of the first 10 minute session. One child in particular seemed to engage with it from the very first exercise she played and one researcher noted after she finished: *"I saw it in your eyes, you went a bit like that and then you moved your arm."* With her it was often possible to see when the feedback was on, not only from her clear arm movements, but also because she would wrinkle her brow slightly in a thoughtful expression or smile. Some of the other children showed when they felt the feedback by smiling or giggling whilst others gave very few signals about what they were feeling but would tell us afterwards *"It buzzed about five times."*

Figure C-1 shows an example of participant A1 playing the piece she were practicing at the time: panel (a) without the feedback and panel (b) with the feedback switched on. Panel (b) shows that the violin is held higher, and in particular, held closer to the target line (at 0) set by the teacher during the calibration for this pupil. The recordings for figures (a) and (b) were made shortly after one another. Notice that in panel (a) the pupil starts at -20cm, whereas for the second play through, in panel (b), she starts at a better position, 0 cm. It was a common theme in her data, as well as some of the other children, that she would begin a piece holding the violin higher if the feedback was switched on, possibly anticipating what the feedback might tell her. Notice the general downward trend in the case without the feedback, a common feature for all the younger children, as they gradually let the violin drop whilst playing. Often they might deliberately do this at one point to get a better look at the strings and then never return the violin to its proper position. In panel (b) she also shows a section of general downward gradient. However, whenever the feedback comes on there is a sharp upwards slope in the graph, showing that she pushed the violin up in response to the feedback. Overall she maintains a position close to the reference position when playing with feedback. The final trailing off is when she brings her violin down at the end of the piece.

It is also apparent that she took much longer to play the piece when the feedback was switched on (25 seconds versus 14 seconds). This is because she played the piece twice in order to make up for an error she had made. We observed that she seemed to make more mistakes when the feedback was switched on, possibly because she was focused on reacting to it.

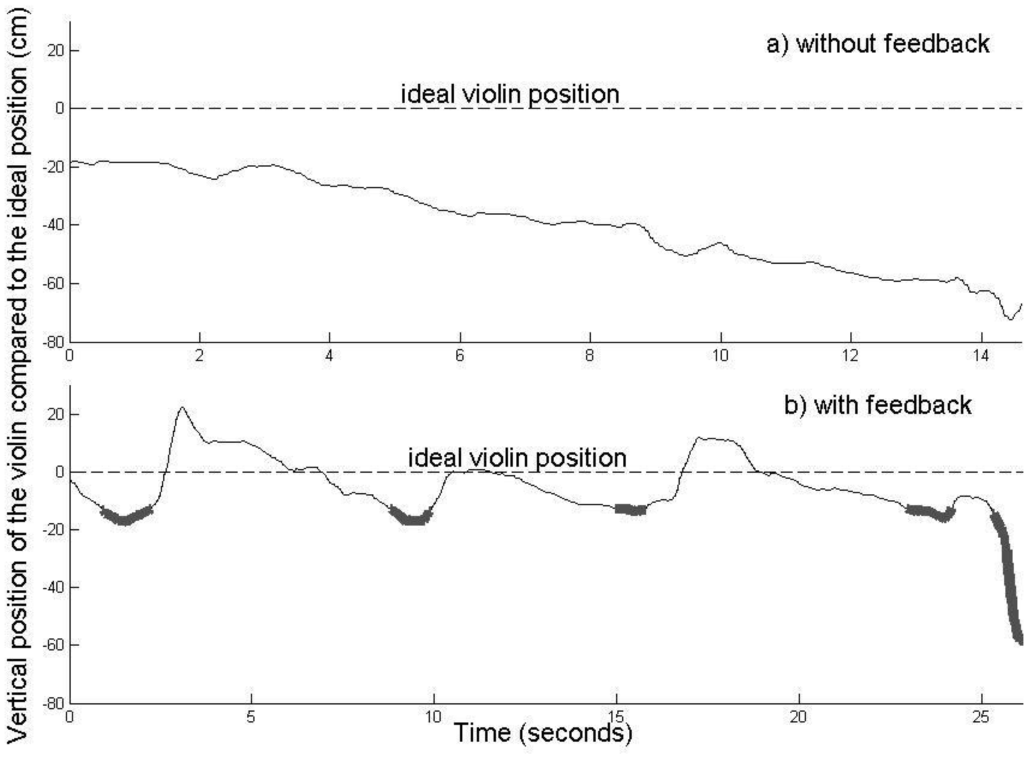

Figure C-1: Panel a) shows the participant A1 playing a piece without any feedback. Panel b) shows participant A1 playing the same piece with feedback designed to encourage her to hold her violin up. Thicker areas indicate feedback is turned on.

The example shown in Figure C-1 is by no means a single case, all the children at different points in their playing responded to the feedback. In the following figures an example is given from each child where there is motion capture data available, unfortunately data is missing from two of the participants. These figures have not been selected to be the best examples of the feedback working effectively, instead they were taken from recordings made in the last session, when all the children played their piece first without feedback and then with feedback in order to give a fair comparison. Figure C-2 shows participant A2 playing much closer to the reference position when playing with the feedback. She has achieved this by spending some time before beginning playing finding the right position to hold her violin where she does not trigger the feedback. As her piece progresses she gradually drops the violin down and triggers the feedback. It is hard to argue one way or the other whether she then responds to the feedback by lifting her violin slightly or whether this is a chance occurrence as there is no sudden lifting like participant A1 did in figure 11. In the previous week when A2 used the feedback she did produce peaks and troughs similar to Figure C-1 showing that she was able to respond to the feedback while playing in some cases.

Figure C-3 shows participant A3 playing close to the reference position with the feedback compared to without the feedback. Like the previous player, the biggest impact comes from starting in the correct position, and a short section of vibration early on helps her to maintain this good start. As she continues playing she gradually lets her violin drop down and triggers the feedback which she does not respond to. When we asked her about the feedback after she finished the piece she reported that it had not gone off since the small buzz at the start. This shows an interesting phenomenon because as the graph shows she was definitely setting off strong feedback over 25 seconds but she was not able to perceive it because she was so focused on playing her piece.

Figure C-4 shows participant A4 making clear large movements in response to the feedback. Her initial measurement shows that she stayed quite close to the reference position without the help of the feedback. With the feedback as well she generally kept her violin position much higher than the reference position and the two times she set off the vibrations she responded to them by lifting the violin much higher. Although we have generally referred to the reference set by the teacher as being the ideal position, it may be better to think of anywhere above the reference position as being an ideal position as one of the experts interviewed in A.3.2 commented, it would be difficult to hold the violin too high.

Interviewing their teacher after the study she commented that she thought this type of feedback had been *“useful”*. In one of the lessons she also mentioned how she *“thought the sound really improved when she suddenly held up the violin.”* This comment was partly addressed to the researchers but also to her student whom she was praising for lifting her violin so well.

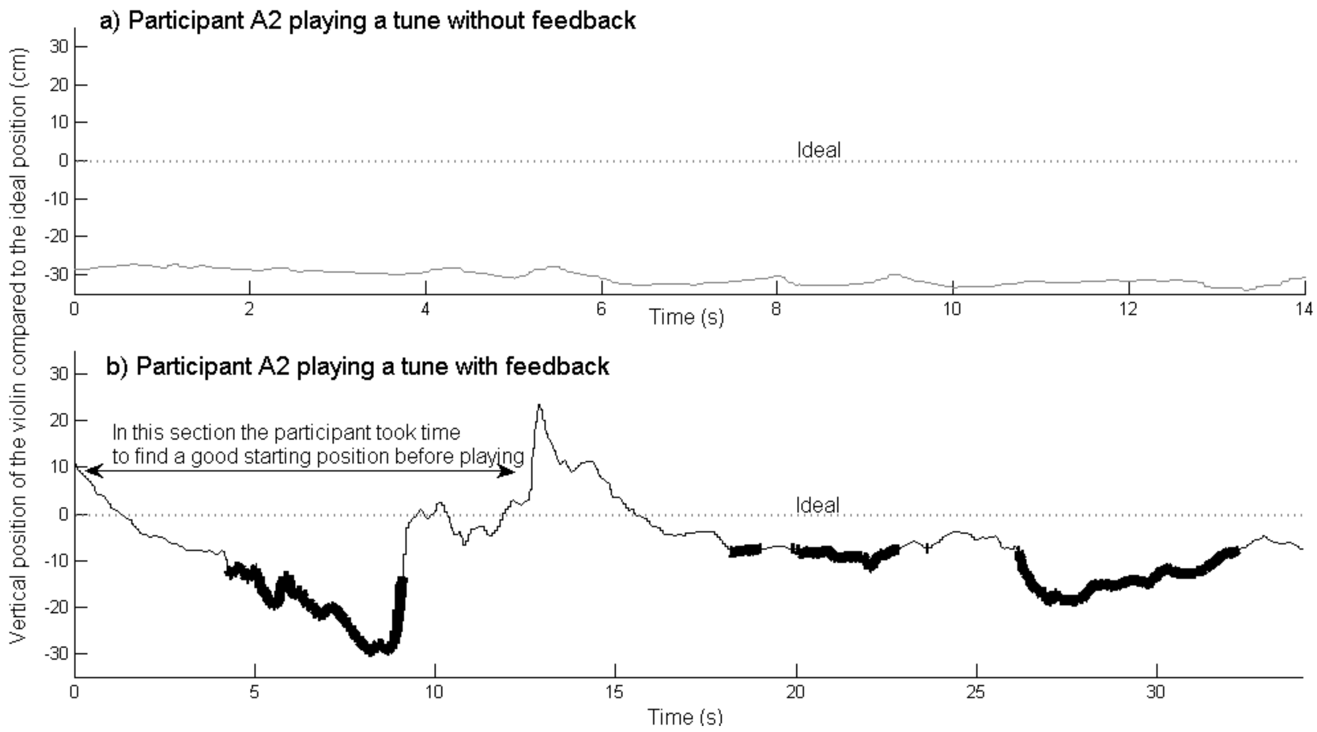

Figure C-2: Panel (a) shows participant A2 playing without feedback. Panel (b) shows participant A2 playing the same piece with feedback. In this case she took time to find her ideal position before commencing playing.

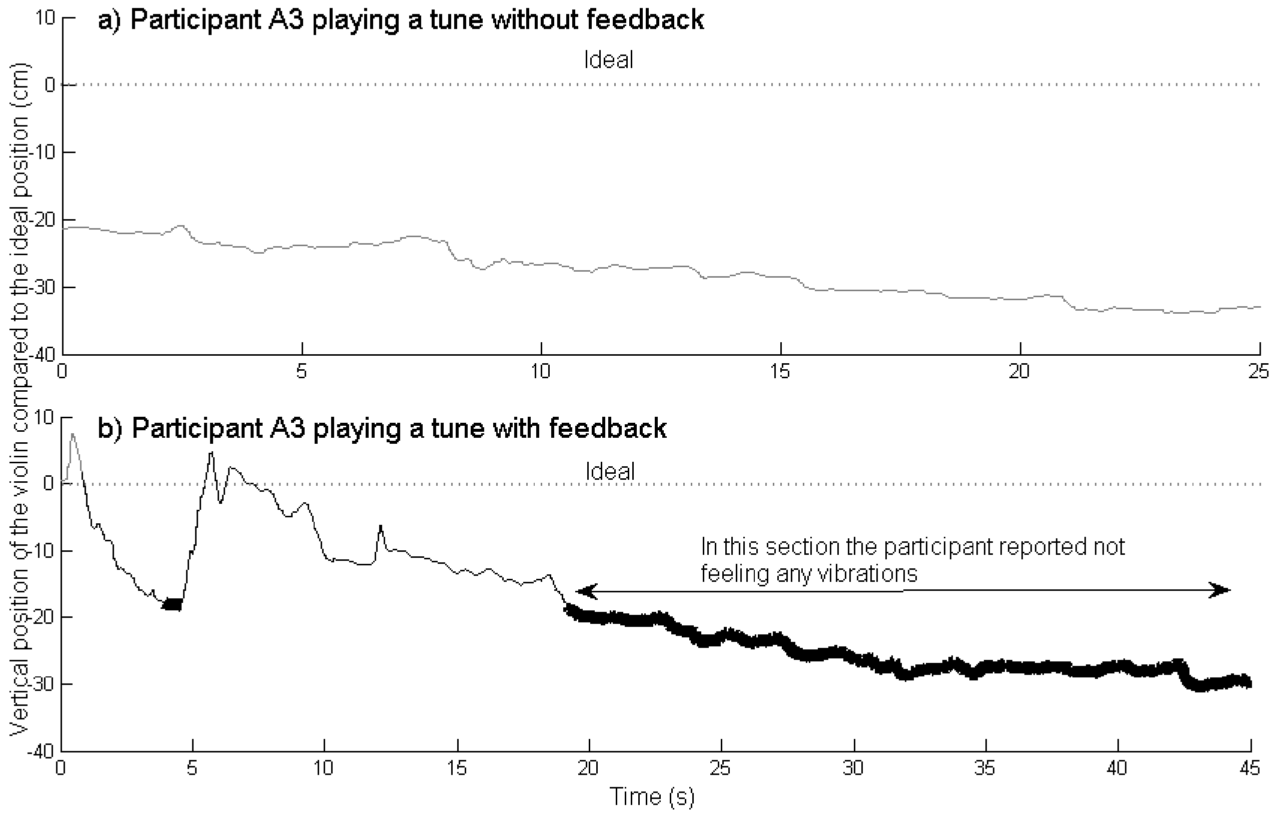

Figure C-3: Panel (a) shows participant A3 playing without feedback. Panel (b) shows participant A3 playing the same tune with feedback. Interestingly she reported after playing that piece that the feedback only came on once at the start, although the graph clearly shows a period of 25 seconds near the end when it was being triggered.

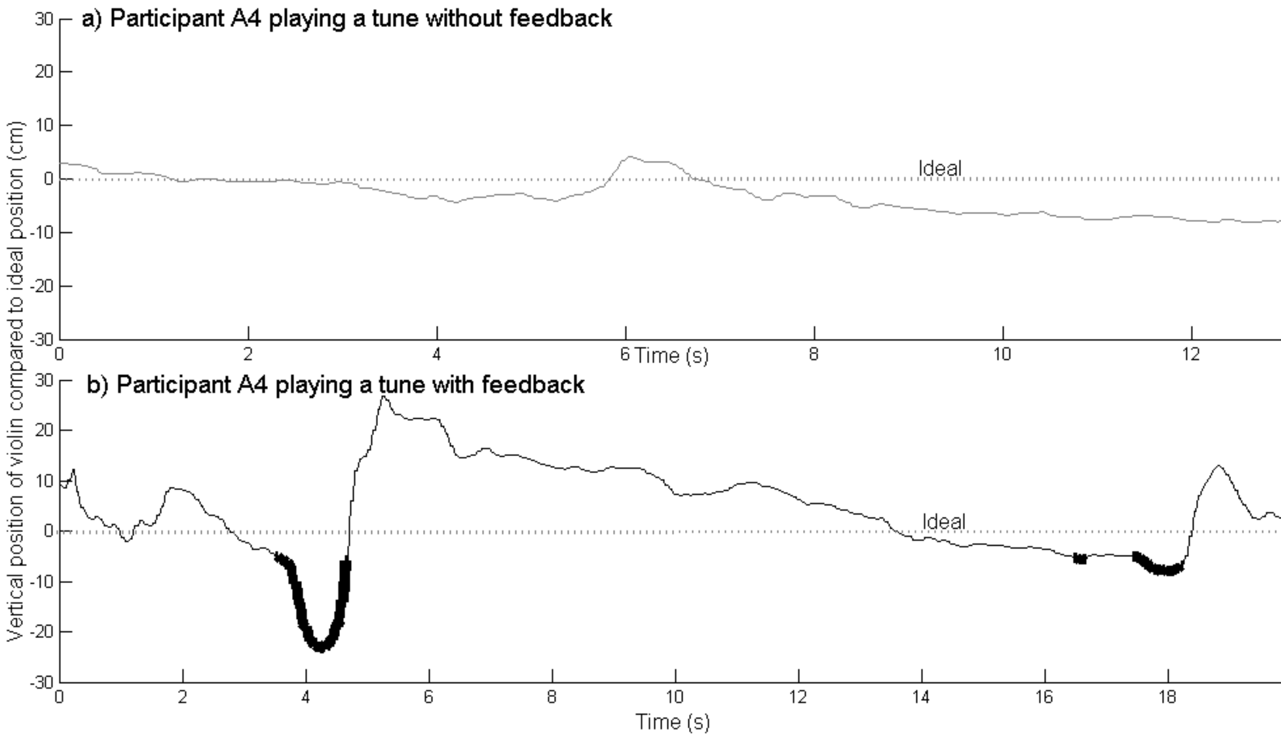

Figure C-4: Panel (a) shows participant A4 playing without feedback. Panel (b) shows participant A4 playing the same tune with feedback. Please note that although the reference position is set by the teacher as ideal, anywhere above the reference trajectory may also be considered ideal as it is very difficult to hold the violin too high.

ii. Straight bowing - Participants B1 and B2

The feedback on straight bowing was more successful with the older children who were at the stage where they now had lots of demands upon their attentional resources: reading or remembering complicated pieces of music; getting their fingers in the correct place; and bowing. These children knew how to bow correctly when this was their main focus but had difficulty doing it while focusing on these other areas. The youngest participant of the group, B1, had particular trouble doing this as she had progressed very quickly to playing more difficult pieces but her bowing had been left behind.

The video data shows how the feedback made B1 glance down at the bow much more whilst she was playing. In her first lesson, without feedback, she only looked down at the bow once during the time she played two pieces. In the following lesson, when she had the vibrotactile feedback, she glanced down at the bow eleven times whilst playing two pieces. The lesson after that, again with the feedback, she looked at the bow fourteen times whilst playing three pieces. These glances are only momentary before she looks back up at the music, but they show that she is much more aware of what her bow is doing. Her teacher was pleased when she noticed this result of the feedback: *"it's quite good to have an eye on your bow and an eye on your music so* *you're swapping back from your music to your bow, music to your bow… You were doing that a bit more I've noticed today."*

The videos show that the result of this increased awareness resulted in more controlled bowing and she was less likely to bow over the fingerboard, which she did in her initial lesson without feedback. This was also observed by her teacher: *"you're not drifting as much on to the fingerboard as you have done. That's been a problem in the past."*

B1 did not express a strong opinion about the feedback verbally. She responded to questions to confirm that she thought it was helpful and then to confirm that she sometimes found it off-putting. As her responses to questioning generally tended to be agreement we also looked carefully at her facial expressions while playing. Similar to many of the pupils in group A, her facial expression changed when she received the feedback. However, the particular change in expression suggests that she had a different kind of experience to the group A pupils. When she plays, with or without feedback, she always has a look of concentration on her face. Without the feedback she will often smile with a self-aware expression when she makes a mistake or gets to a certain part of a piece such as the end of a phrase. When playing with the feedback something slightly different happens, her face moves much more but by smaller amounts, alternating between a slight smile and lips pursed together. Her general demeanour when the feedback is switched on also seems to be tenser, and her playing seems to flow less. This seems to indicate that she is finding the task more demanding when the feedback is switched on. When one of the researchers mentioned to her that she was frowning while playing she answered that it was because *"I couldn't remember the notes".* This confirms that she was finding playing that piece with the feedback demanding.

Although B1 only had three sessions with the feedback her teacher used the study as a starting point for working on her bowing during the months following the sessions. Significantly, the teacher and student talk about bowing movements in terms of vibrotactile feedback.

"We've talked about what we did – how she felt the feedback when she was going too far forward… and how she had to remember to counteract that by bringing her arm back, and she's got to try and do that herself and keep an eye on her bow… rather than just ploughing through she's got to make conscious decisions about where she's bowing and how she's bowing."

Participant B2 also needed to bow straighter. Like participant B1 it was observed that she began to glance more at the bow when she had the feedback. The motion capture data also shows her being guided to a better bowing trajectory by the feedback. Figure C-5 shows a good example of this guidance effect – initially her bowing trajectory was too close to the fingerboard, but after the feedback was switch on she immediately adjusted her bowing to a better position. We have not included figures from participant B1 because her more gradual response to the feedback cannot easily be displayed graphically.

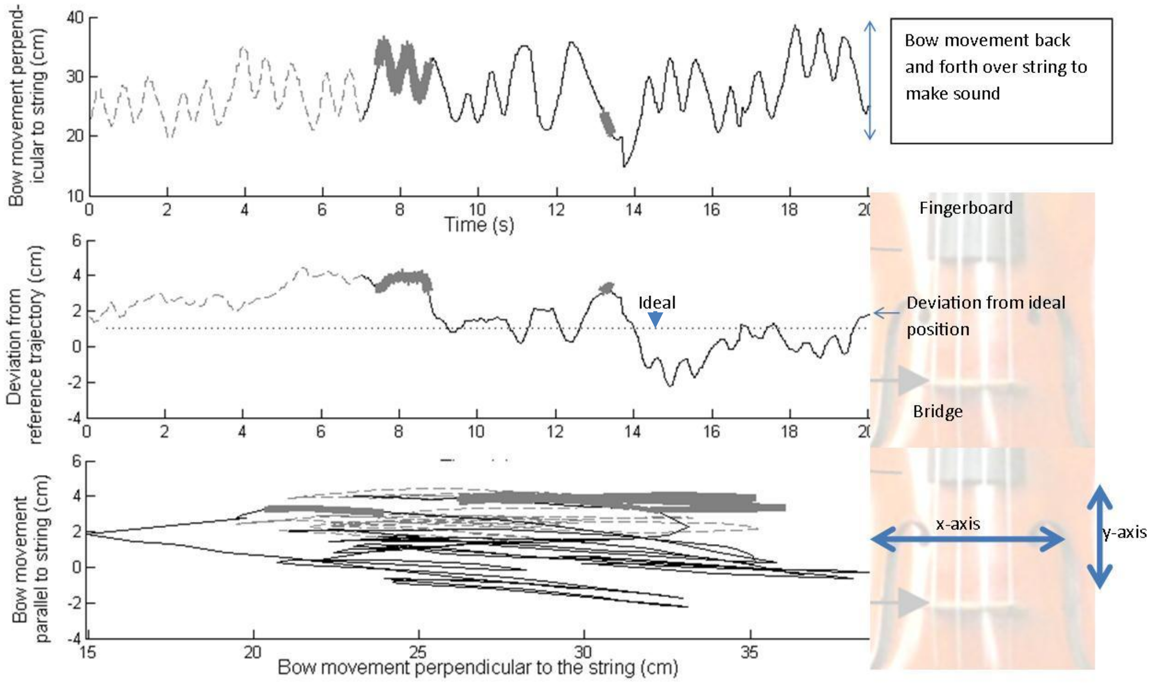

Figure C-5: Participant B2 responding to feeling the feedback. Top panel shows the movement of the bow hand perpendicular to the string over time. The middle panel shows the movement of the bow hand parallel to the string over time. The ideal position is 0; this would mean the participant was playing on the reference trajectory. The bottom panel show the movement of the bow hand over the whole playing time as if viewed from above. Bold areas indicate when the participant received vibrotactile feedback.

iii. Using more bow - Participant B3

As previously mentioned, the MusicJacket was adapted to encourage long bows in order to meet the needs of a particular student in group B, B3. This student's main difficulty with bowing was that she played with a small amount of the bow. Figure C-6a shows how she regularly used approximately only 10cm of bow (a typical violin bow is 75cm long). During the study her teacher spent time with her encouraging her to use more bow. The result of this is Figure C-6b which shows how she had progressed by the end of the lesson to using 36cm on the longer notes. However she often reverted back to her old style of bowing by the next lesson.

When the pupil tried the feedback to improve straight bowing she did not get much feedback. This is no surprise as it is easier to bow straight when using very little of the bow. Her teacher pushed her to use more bow by repeatedly reminding her whilst she was playing. However this additional focus upon bow length seemed to prevent her from noticing the feedback on straight bowing, for afterwards this pupil reported she did not feel any vibrational feedback even though the system recorded that it had triggered.

In session five she was given feedback to encourage long bowing. Figure C-6c is a result from that feedback session, in which she began to regularly use over 50cm of bow. The effect occurred almost as soon as she tried the feedback making it quicker than the lesson which produced Figure C-6b. Clearly this was a breakthrough moment, and both pupil and teacher were very pleased about the new approach. In an interview after the study her teacher commented *"it was very useful and I think that it worked really well"*.

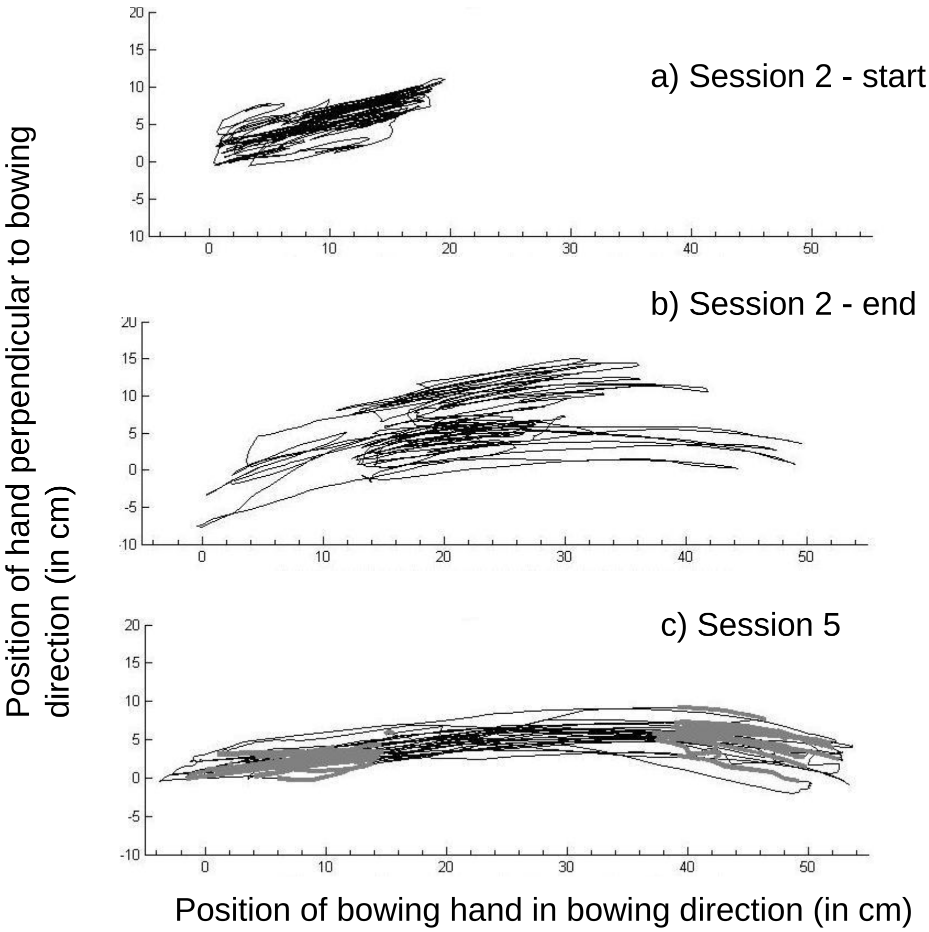

Figure C-6: Progression in bow length for participant B3 shown over several sessions. Panel a) Bowing at start of session 2, without vibrotactile (VB) feedback. Max bow length = 12.5 cm; Panel b) Bowing at end of session 2 which had concentrated on increasing bow length using normal teaching methods. Max bow length is 36.7 cm. Panel c) bowing at end of session 5 with VB feedback designed to encourage to use more bow. Max bow length = 55.3 cm.

iv. Calibrating to personalise feedback – Teacher B

Calibrating the MusicJacket with a good reference stroke was essential to using the feedback. At the design stage calibration had been viewed as a necessary functional process where the participant and their teacher carried out a series of processes as instructed by the researcher to fulfil a need of the technology. However during the study we found that calibration can be much more than this; it allows the student and teacher to specify where they wanted the feedback to be triggered. In the fourth session, in particular, a lively discussion ensued around this process of calibration. It began with the teacher noticing that her pupil, B2, was playing too close to the bridge of the violin. After talking to the pupil about it she ascertained that the feedback was not being triggered when her pupil bowed close to the bridge and she brought this to the researchers' attention by saying "*Yeah it should really have been going off at that point, I think, and it wasn't.*"

The teacher and researchers together experimented with several different bow positions to find where the feedback was or was not triggered and the teacher pointed out when the feedback ought to be coming on. For example the teacher asked her pupil *"Now put it closer to the bridge because you were playing about here weren't you. Right is that buzzing now?"* and then on finding out that the feedback was not being triggered near the bridge she asked the researcher to adjust MusicJacket "*So we need a way to get that buzzing there. Because if you play there on the E string it's going to sound squeaky!"*

In response to this, the researcher tried making the margin for pupil error around the reference stroke smaller so that if the pupil only moved their hand off the reference path by a small amount the vibrations would be triggered. However, this resulted in making the area near the fingerboard too strict for her bowing style while not making a big enough difference at the bridge. Next, the reference stroke was recalibrated to ensure the reference stroke was in the correct position. However, the problem remained that even on a strict setting the feedback was not sensitive enough for the needs of this particular pupil around the bridge and simply making the feedback stricter would have inhibited her freedom of movement too much.

Finally, the reference stroke was recalibrated again but positioned so that the stroke was closer to the fingerboard than the bridge. This reference stroke was not the perfect position to play in but it meant that the vibrations were triggered more easily near to the bridge than at the fingerboard as the pupil margin for error each side of the reference stroke is symmetric. The final calibration gave feedback where the pupil needed it and caused her to play her next piece with her bow well positioned. However after playing the piece for a minute she chose to stop as she found the strict feedback uncomfortable. This shows that while it is important to be able to set feedback which challenges an individual's common mistakes, pupils also need to be able to adjust feedback for comfort as well as for learning, especially as they get tired near the end of a lesson as in this case.

This 'misuse' of the calibration of the reference stroke enabled this pupil who had a particular *"bridge problem"* to get feedback that met her specific needs. Moreover, in the course of the discussion that led to this calibration, the teacher also learned more about how the system worked. It caused her to ask questions about what the system was able to measure and how it did it; and what this meant when using the system – that it easier to give feedback on poor bow angle than a straight bow in a bad position. It also made her clearly state where the bow should and should not be positioned, effectively setting a goal for her pupil. The calibration was then used to implement this goal.

C.5 Discussion

The findings from this study show that real-time vibrotactile feedback can be a helpful tool for teaching violin lessons and can result in improvements in playing. The exact ways it does this is strongly dependent upon the individual learner and their teacher. The feedback had to be personalised towards the learning goals of that the teacher had set for the student to be useful. In cases where it was not then it was ineffective and in some cases may have been overloading, for example when participant B3 was getting feedback about straight bowing while trying to focus on long bowing. Personalising the feedback towards learning goals took two forms: first, the teachers chose the aspect of feedback that was most relevant; second, one of the teachers then went on to specify exactly where she wanted her student to feel the feedback in order to tailor it to that student's particular problem with bowing. The experience of using the MusicJacket in the classroom suggests that it might be useful to reframe our thinking about calibration so that instead of recording the ideal bow stroke we aim to facilitate people in specifying where they would like feedback to be triggered.

As well as having feedback which is suited to the learner's goals we also found that it is important that feedback is used on suitable tasks. A person only has a limited amount of attention available and using vibrotactile feedback can interfere with playing and vice versa. This means that playing with vibrotactile feedback can make a piece more difficult and playing a difficult piece can prevent people from noticing the feedback. In practical terms this means that feedback should only be used on easy and intermediate level tasks and that care should be taken to only give feedback on one aspect of playing at a time. Task must not be too easy however as this can mean they lack challenge and interest for learners. Again whether a task is easy or difficult is down to the abilities of the individual learner. This is one of the reasons why the teacher's input was so important because they had an intuitive understanding of their pupils and their strengths and weaknesses and could choose just the right exercises and pieces for them to learn with the feedback.

In addition to its direct effect on playing the vibrotactile feedback was also helpful in another way: it promoted discussion between teacher and pupil and gave a different way of talking about the body. This was particularly apparent in the group B, the older children, where in the interview the teacher said she still referred to their experiences using the feedback some weeks after the study had ended. This is an interesting and unexpected finding; before the study it had been feared that feedback could come between teacher and pupil because it could only be felt by one party, but instead it enabled discussion because asking whether the feedback had come on became a new line of questioning, which could then initiate a dialogue about movement. This raises questions about whether this promotion of dialogue about the body is only a property of vibrotactile feedback or whether feedback in another modality such as ambient visual feedback would also elicit the same kinds of discussion. Vibrotactile feedback is felt on body and therefore may be more likely to promote thinking about the body; however ambient visual feedback is shared with everyone present which may facilitate discussion.

The episode of the feedback to encourage long bowing demonstrates a different configuration for feedback, namely the vibrations became a positive indicator to be aimed for whereas previously the vibrations pointed out mistakes. Initially we had been concerned that having the feedback as a negative indicator may not be motivating and may leave participants feeling that they have been 'told off'. Certainly the added tension we saw in some of the students from group B could be attributed to that. However, we have also seen cases, (mainly in the studies with adults) where participants expressed pleasure at managing to play without triggering feedback showing that negative feedback can still be rewarding. Our rationale for choosing the vibrotactile feedback as a negative indicator was that the sensation of the vibrations was somewhat annoying and would not be suitable for a reward. However the case of the long bowing application shows that a vibration can be positive so long as it is short. Nonetheless, it could be argued that a different modality such as twinkling lights might be even better suited for giving a positive reward.

The teachers also pointed out the long term success of real-time feedback will be dependent upon frequent and sustained use. In particular, feedback which could be used in practice at home to reinforce progress made in lesson time would be a promising area for research. This raises new design challenges and research questions. In terms of design, producing something which is inexpensive, robust and simple enough to use at home will require a different design approach to the one used for the MusicJacket, which is reliant upon a costly and complex motion capture suit connected to a laptop. In terms of research this raises questions about how people will choose to use feedback once they are given full control of it and are able to take it into their own home and whether they will feel motivated to use it at all. This thesis addresses these challenges and questions.

Appendix D – Interview Plans and Questionnaires

D.1 Interview Plan from Chapter 5

This interview was semi-structured so questions were sometimes phrased differently and follow up questions were asked based on the participant's initial responses.

1. Please put the feedback in order of preference?
    - Paper slips with the names of each prototype are provided to participants for them to arrange physically
    - Now go through each prototype in turn and ask why they put it in that position
    - Follow up with further questioning
- Could you imagine using technology like this in your own practice at home?
    - Follow up with further questioning
- Imagine I was making a device just for you. What would you like it to measure, and how would you like it to show it to you?
    - Follow up with further questioning

D.2 Questionnaire from Chapter 6

(Please note layout has been condensed to save space)

Name:

Why did you choose to come to the summer school?

Which type of feedback did you prefer the visual feedback or the vibration feedback?

Why was this?

If you had one of these systems to use at home how do you think you would you use it?

(For example: Would you use it for particular exercises or scales or in ensembles…?)

Vibrotactile feedback

On average how often do you think you were aware of the **vibrotactile feedback** when it went off while you were playing? (please underline)

- less than 10% of the time (I almost never noticed it)
- between 10% and 30% of the time
- between 30% and 50% of the time
- between 50% and 70% of the time
- between 70% and 90% of the time
- over 90% of the time (I almost always noticed it)

Please give any comments you can about the vibrotactile feedback:
(For example: In what situations is it more noticeable? Is it comfortable? Do you mind using it in a group?)

Visual Feedback

On average how often do you think you were aware of the **visual feedback** while you were playing? (please underline)

- less than 10% of the time (I almost never noticed it)
- between 10% and 30% of the time
- between 30% and 50% of the time
- between 50% and 70% of the time
- between 70% and 90% of the time
- over 90% of the time (I almost always noticed it)

Please give any comments you can about the visual feedback:
(For example: In what situations is it more noticeable? Is it off-putting? Do you mind other people being able to see it?)

Practice strategies

On average how many hours a week do you practice?

If you are stuck on part of a piece or an aspect of technique how do you go about improving? (If possible think back to the last time you were stuck and use it as an example)

Thank you

Thank you for taking part in my study and answering this questionnaire. Please feel free to give any other comments you would like to here:

D.3 Interview Plan from Chapter 7

Section 1 – Home practice

When you normally practice at home what do you do?

- How often are you able to do this?
- Do you have to be reminded?

When you had the prototypes were you able to fit them into your normal practice?

- How difficult was this?
- Did you have to change what you usually do?
- Did you mind using them?

Section 2 – Comparing modalities

Which of the two types of feedback did you prefer?

- Why was this?
- Was this the same for the school practices as for home practice?

For the following questions I'm going to use a smiley face rating scale similar to Janet Read's. See the end of the interview for the question sheet we'll go through together.

Can you rate the lights and the vibrotactile feedback in terms how easier or difficult they were to notice while you were playing? (Why?)

How about understanding what they were telling you to do – how easy or difficult did you find it to understand the two types of feedback? (Why?)

And comfort – how would you rate each type of feedback in terms of how comfortable you felt using it? (Why?)

How well do you think each type of feedback worked for improving your posture? (Why?)

How would you rate each type of feedback in terms of enjoyment compared to normal practice – did you find it added fun or made it more boring? (Why?)

Overall, how would you rate each type of feedback?

Section 3 – Looking for signs of preferred modalities

Attitudes about reading music

When you are playing music do you prefer playing it from the notes or off-by heart?

When you're learning a piece of music can you learn it just from the written notes or do you like to hear it played first?

Do you prefer sight-reading or aural tests?

How easy or difficult do you find it to watch Mr Philips (the conductor) while playing in Sinfonia? *(Use rating scale here)*

Thinking about the body when playing

When your teacher wants you to change your posture or how you are holding the viola or bow what does she do? E.g. does she physically move you, does she describe the movement…

- Can you give me an example?
- How easy or difficult do you find this to respond to? *(Use rating scale here)*
- Is there another way you would prefer her to show you?

If I ask you to play your piece holding your violin up all the way through – what do you do/imagine to achieve this?

- How do you think about your arm or viola while you are playing?
- How easy or difficult do you find this task? *(Use rating scale here)*

When I asked you to visualised your head as a ping-pong ball on top of a fountain at the beginning of the sessions, how easy or difficult did you find doing this? *(Use rating scale here)*

- Did you think it changed the way you were standing?
- Would you prefer it if I had moved you physically?

Other potential areas which may show preferences for modalities

What are your favourite school subjects and why?

Section 4 – Design ideas

Imagine I was design a new prototype that could measure how straight you were bowing, how should I design the feedback?

- How would it tell you?
- What would it look/feel/sound like?
- Where would you use it – at home, in school, in lessons…

(I'll have paper and coloured pens for drawing)

Question Sheet

Rating the feedback

Question	Rating
Noticing the lights	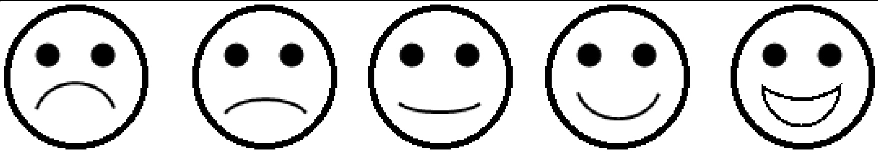 Very difficult Very easy
Noticing the vibrations	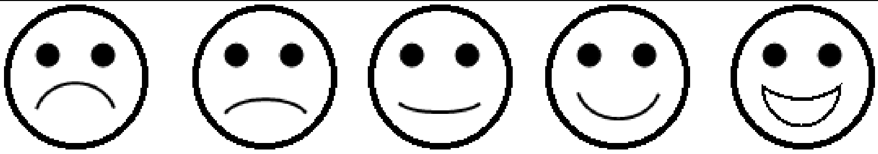 Very difficult Very easy
Understanding the lights	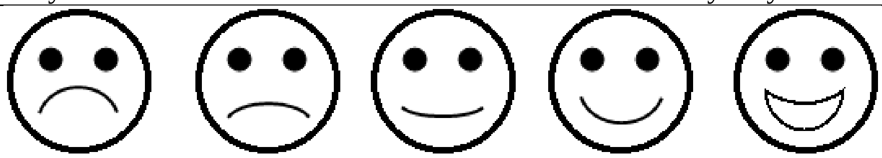 Very difficult Very easy
Understanding the vibrations	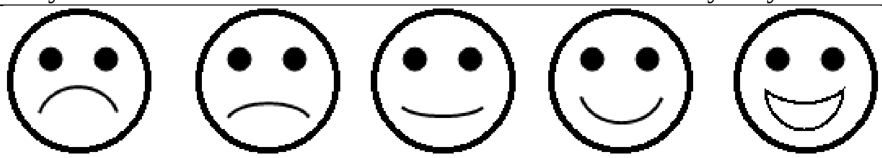 Very difficult Very easy
How comfortable did you feel using the lights?	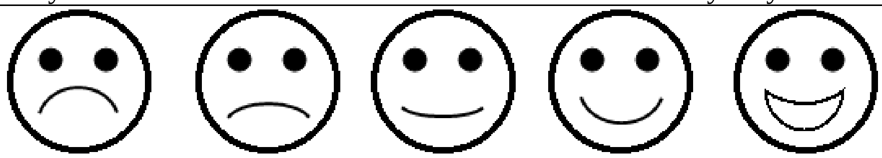 Very uncomfortable Very comfortable
How comfortable did you feel using the vibrations?	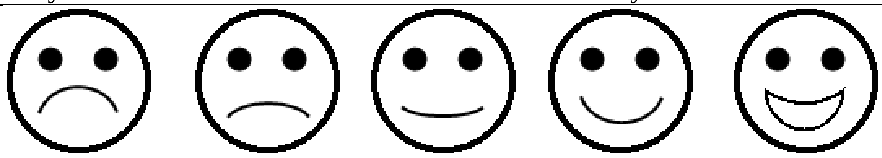 Very uncomfortable Very comfortable
How well do you think the lights worked for improving your foot position?	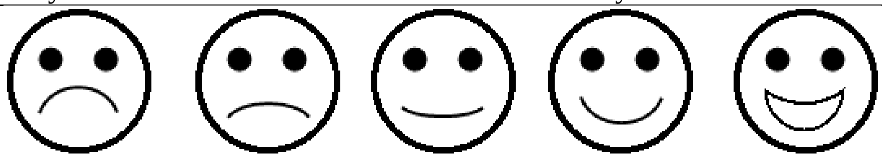 Very badly Very well
How well do you think the vibrations worked for improving you violin position?	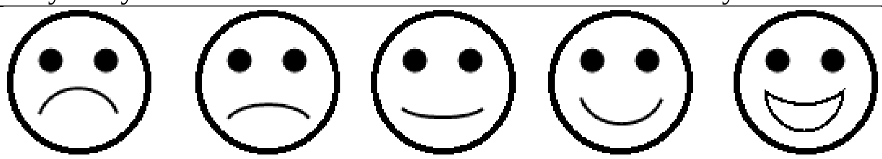 Very badly Very well
How much did you enjoy using the lights?	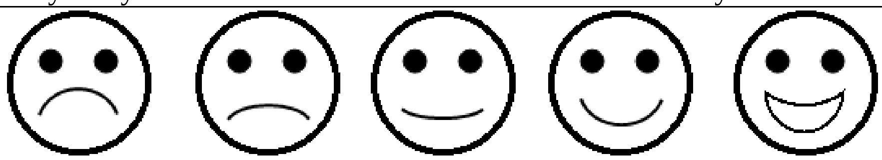 Not at all Enjoyed a lot

How much did you enjoy using the vibrations?	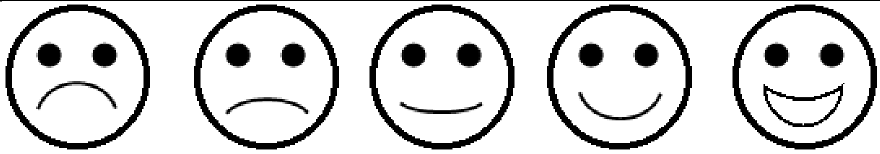 Not at all … Enjoyed a lot
Overall how would you rate using the lights?	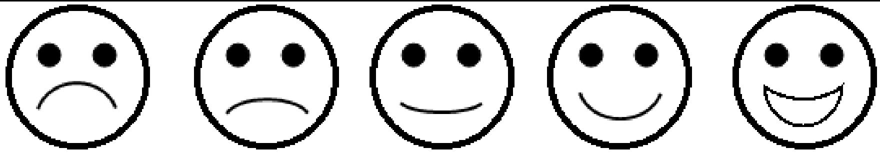 Very bad … Very good
Overall how would you rate using the vibrations?	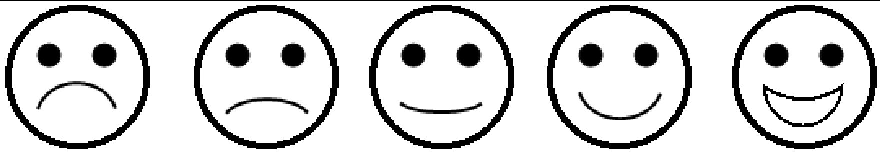 Very bad … Very good

Other questions

Watching the conductor	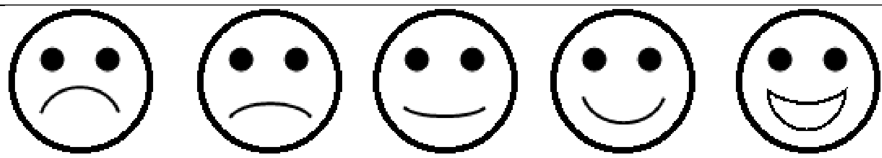 Very Difficult … Very Easy
Responding to instructions about posture from teacher	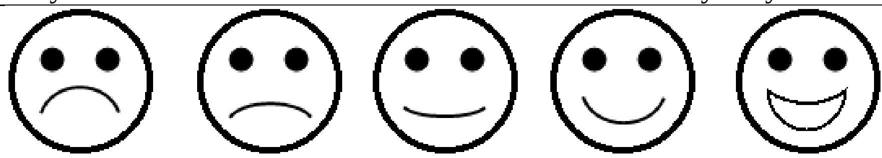 Very Difficult … Very Easy
Focusing on holding up violin while playing	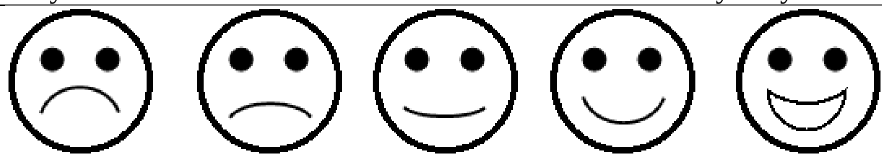 Very Difficult … Very Easy
Visualising the ping-pong ball on the fountain to stand up straight	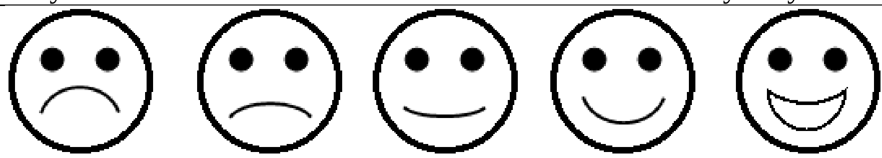 Very Difficult … Very Easy

Appendix E – Analysis of Individual Differences in Chapter 7

Alice

Alice's positive personality showed through particularly in her response to questions about how she would design a new feedback system to help people with straight bowing. Her system used negative vibrotactile feedback combined with a rating scale which could give positive feedback. This shows how multimodality can be used to give a more balanced style of feedback which encourages players to focus more on their achievements than their mistakes. Her first suggestion was *"a bow that dispenses sweets when you play right"* her reasoning was:

"Because that means it's focusing on what you're doing well, so it's making you do it well rather than going 'don't do it wrong'"

She then went on to clarify that it wasn't that she wanted to be *"rewarded"* but just that she wanted it to focus on what was being done right rather than preventing mistakes.

Alongside this wish for a positive focus she also liked the vibrators or "the fuzzy things" as she called them. She considered the idea that they could be used for positive feedback but felt *"they might be buzzing quite a lot*." In discussion with the interviewer she finally fixed upon the idea of having the vibrators to point out mistakes and then having additional feedback at the end giving a rating of the piece as a whole. The additional feedback she describes is given as a percentage *"like you did 60% with the bow straight and 40% with it not straight."* When asked if the thought of getting her percentage higher might make her more inclined to practice she replied yes.

This design reflects elements of her personality which also came out in the sessions; the main ones being a desire for a positive outlook and recognition of achievement and improvement. For example she would point out in a session when she had done something well, whereas some of the others were more likely to focus on mistakes. She would ask questions which asked the researcher how well everyone was doing. She also suggested games with the lights which offered a chance to show achievement. Her description of how her teacher would correct a mistake also shows that he balances a negative comment with a positive recognition of overall achievement:

"I'd got to three bars and then he stopped me and said: 'wait a minute, it's brilliant, but can you make your legs a bit better.'"

Alice's design tries to strike a positive balance between praise and constructive criticism in the same way her teacher does in a lesson.

Notice that in her design she chose to use vibrotactile feedback even though it would have to be set-up as negative feedback rather than real-time visual feedback which could more easily be set to give positive feedback. Of all the students she showed the strongest preference for one feedback type over another. She described the lights as *"distracting"* and she *"didn't want to look at the lights instead of the music."* This difficulty with splitting her visual attention between the music and the lights was also apparent in the videos of the group sessions in which she would sometimes play for a whole piece without noticing that the lights had come on. The importance to her of looking at the music was also echoed in her description of trying to look at the conductor which she felt was *"quite hard because I'm really watching the music because it's quite hard notes."* In her description of how she would normally go about improving an aspect of playing (in this case we talked about bending the fingers of the left hand correctly as this is something she felt she needed to improve) she said she would begin by learning her piece by memory so that she didn't have to look at the music and could look at her fingers while playing. This suggests that she prefers to have a single visual focus where possible and tries to remove visual distractions. Moreover, the videos of her home practice show that when playing her pieces she glances down at the violin frequently either to check her fingers or her bow showing that she already has to split her visual attention without the addition of the lights as well. She did not feel the vibrations were distracting in the same way saying *"you don't even look at them, you just feel it,"* and later in the interview adding *"you don't have to focus on them."*

Although she preferred the Buzzy Jacket, she felt that the lights were more suited to the group practices because they publically showed how each person was doing:

"because you could kind of see how much other people's lights were going on and you could kind of…"

This demonstrates that peer awareness is part of the power of visual feedback as already described in the observations of the group sessions.

In terms of learning focus she felt that the jacket was more useful in her practice because keeping her feet still was something that was easier to improve:

"I think they worked well, but I don't think it took very long. It was almost as if once I knew it was there, I stopped crossing my legs anyway."

This is a perceptive comment because we found that even the presence of a piece of paper which the children had drawn their feet on was effective for preventing the children from moving their feet while playing. On the other hand she felt that the Buzzy Jacket gave her feedback on something that she would have more difficulty improving without help:

"Yes, that worked really well, because I wouldn't normally notice."

Learning quickly with the Footboard and Twinkly Lights does not discount its usefulness. Instead it shows that it can play a valuable role within a programme of feedback, but that it should only be used for a short period before feedback is given on a different aspect of technique. This shows how a difficult subject in technique like posture can be broken up in to smaller elements which are more easily worked upon and then gradually built these up over time through a programme of feedback. This can prevent the cognitive overload that could come from introducing too much feedback at once.

Bea

For Bea enjoying practice is a priority, perhaps more so than simply finding the system that works best at pushing her into the right position or movement. In her design for feedback she chose to use lights despite stating earlier that the vibrotactile feedback was more effective. Her example shows that when designing feedback for the long term, short-term effectiveness may not be the chief consideration, enjoyment is also important especially if the learner less motivated by achievement goals and more by the pleasure of playing.

Bea found the jacket most helpful because she felt in had the most relevant learning focus for her. When asked about her preferred modality of feedback Bea chose the lights:

"I did like the lights. They were quite cool. And sometimes you couldn't really feel the vibrations."

However, this preference is not shown very strongly on the smileyometer. She rated the vibrotactile feedback more highly in terms of comfort, understanding and effectiveness. She found the lights difficult to understand because the feedback didn't give enough detail of how to improve:

"Because they just all came on, and you didn't know which foot or which part of your foot it was, so that was quite hard."

Whereas she found the vibrotactile feedback being on the body made it easier to interpret:

"That was very easy because you just had to feel where the vibration was and stuff, and what it wanted you to do."

The reason she gives for rating the vibrations as being more effective than the lights is that the vibrations forced her to change her playing:

"It really helped you get it higher, because if… Yes you're just wanting the vibration to sort of go away, so you just have to change it."

The only aspect where she rated the lights more highly than the vibrations was in terms of how much she enjoyed using them:

"It was quite fun, because I liked the way that there were different patterns in the lights and stuff."

Despite all the good things she had pointed out about the vibrations, this was enough for her to choose to use lights in her design for a new feedback system. She describes a system which lights up when the bow goes off the correct path with colour used to communicate exactly what the mistake is. When describing the lights in her design she views noticing and understanding them as being less problematic than she had described the lights with the footboard:

"I would quite like to have lights because they're quite easy to just tell […] because if you're looking at the music stand and you just see them and then you'll kind of know."

This may be because she feels that set up in the right way, with the additional detail given through the colour, her lights would be easier to use than those that came with the footboard. Overall, however, we can see here that she is choosing the feedback that she most enjoyed using perhaps with the idea that she can make it more effective by tweaking the design rather than choosing the modality of feedback which she considered worked best in the study.

Enjoyment is an important consideration when designing a long term programme of multimodal feedback. While music is widely acknowledged to require a great deal of dedication it is also vital for it to be enjoyable, particularly in formative years. There are times in a student's learning where they may need to be pushed and feedback which is perceived as being more forceful such as vibrotactile feedback may be very useful then. At other times it is important for a student to have the opportunity to enjoy playing, especially for a student like Bea who did not seem so driven by other motivators such as achievement goals or a desire to perform. For Bea the feedback which made playing more enjoyable was the lights.

Carl

Like Bea, Carl also perceived the vibrations as being more forceful in the way they encouraged him to correct mistakes. However he felt this was an asset, most likely because he felt it would make him improve faster. He was motivated strongly by a desire to perform and show others his musical abilities. His reason for practicing was mainly to improve so that he could perform. This approach to practice would lead to choosing the type of feedback perceived as being most likely to improve his playing.

Carl found the footboard more useful than the jacket because the learning focus was more relevant: he felt that he had particular difficulty keeping his feet still and less trouble with upper body posture. When asked which modality of feedback he preferred he chose the vibrations:

"The vibrations because that ***made*** *[his emphasis] you go up; the lights you could just not notice them."*

Despite making this point that the lights were hard to notice he then rated the vibrations as harder to notice on the smileyometer saying:

"Sometimes they were quite hard to feel if you weren't bent as much, like your back - *it wouldn't be as powerful."*

He also found it difficult to distinguish which motor was vibrating and for this reason rated the vibrations as more difficult to understand. Therefore rather than interpreting his reason for choosing the vibrations as being to do with awareness it is important to look at his phrasing. He emphasises that the vibrations "***made***" him lift up whereas *"you could just not notice"* the lights. Through this he is indicating that the vibrations are more forceful in promoting a good position than the lights which you could choose to ignore. So here it is not so much a case of whether a player was able to notice them or understand them (which are the questions on the smileyometer) as how strongly the feedback promotes the correct behaviour. Carl appreciated the way the vibrations *"made"* him do the correct thing. It is interesting here to note that although he found the vibrations forceful, he also describes not noticing them, showing that forcefulness is not simply how effective feedback is at attracting attention.

He also chose to use vibrotactile feedback in the design part of the interview. His appreciation of the forcefulness of the vibrotactile feedback is in line with his approach to practice. He liked to perform and get the appreciation of his peers. Before group sessions he would play the piano to the others. After using the video camera the first week he played the recordings of himself to the others in the group. He took part in many out of school concerts and also performed in drama groups. His approach to practice seemed to be functional; he did it mainly to be able to perform better. For example he describes practicing only twice week normally which is quite low compared to the others participants but in the run up to an exam or concert he practices every day. From this functional perspective on practice the forcefulness of the vibrotactile feedback is an advantageous quality that will make it effective. This contrasts with Bea's approach which was to choose the feedback she most enjoyed.

It is also interesting to comment that while he was happy for the vibrotactile feedback to make him do things but he was very clear that he did not like his teacher touching him or physically moving him to correct things, he much preferred verbal descriptions or demonstrations as teaching methods.

It is important to note that his responses in the interview are inconsistent. For example, he said he preferred vibrotactile feedback as a modality but then rated it lower on the smileyometer. This is a methodological issue since the questions in that section refer to how effective the vibrations or lights are at improving posture or foot position (respectively). As such it is strongly connected to the element of technique the feedback was applied to. Since he found the footboard more relevant to his learning this may account for his lower ratings for the vibrotactile feedback despite it being his preferred modality.

These contradictory responses may also be explained in part by his ideas about how he thought the system works. His description shows that he had the idea that sensors and actuators would have to be in the same place. For example when talking about using lights as feedback for bowing in the design exercise he worked on the assumption that the feedback was built into the sensing system:

"Not for the lights, like if it was a sensor from the shoulder because I may not be able to see it."

Even with the interviewer suggesting they could be wireless so that they could go anywhere his responses still show concern for the practicalities of how the system would function:

"Yes, or it could be electrical pulses where you get it through the bow or something and you have a sensor here and you feel it when it's made contact with the fingerboard."

Therefore when he says he would choose vibrations for the "particular" case of improving bowing it maybe in some part out of concern for how the feedback would be connected to the sensors. This strong connection in his mind what is being sensed and how feedback is given is likely to lead to inconsistencies in his answer to questions such as those on the smileyometer which try to separate the two.

To summarise, Carl shows an example of a functional approach to practice. He does not show a strong preference for a particular modality, but does appreciate the forceful nature of the vibrotactile feedback which fits well with his approach to practice. When designing a programme of feedback for a student it is important to understand how they practice. For someone like Carl who is driven by external goals such as performances and exams a more functional approach to feedback can be taken. His periodic practise regime which is based around these external goals may also particularly favour feedback which is effective with short-term use.

Daphne

Daphne had a laid back approach to music practice, which she mainly seemed to do for enjoyment. This relaxed approach also applied to her opinions about the feedback where she showed an appreciation for the usefulness of both the Buzzy jacket and the Twinkly Lights and did not wish to choose one over the other. Her own feedback design at first seemed at odds with her fun-loving personality because it not on the surface a particularly fun design. However viewed from a different angle it shows that she does not need external rewards or overtly game-like elements to make practice interesting – for her getting feedback on playing can be interesting in itself.

Daphne was a cheerful student. In the group sessions and interview she laughs a lot and makes jokes. She describes enthusiastically how her teacher "kind of makes it funny" when reminding her of bad habits and says this makes it easier for her to remember. It is also interesting to point out that she enjoys this despite it being a criticism of her playing. This is a contrast to one of the other students with the same teacher who reported finding it very hard to accept criticism at first. Her description of her practice routine shows that she is not highly driven to practice but when she does do it, it is intrinsically motivated:

"I sort of play it when I feel like it or when I have time and I used to have to be nagged but seeing as I haven't played the violin for ages I'm starting to do it without being nagged now."

She also describes leaving her violin out so that she can pick it up whenever she likes.

She was also inquiring and interested to see things from different perspectives. For example, during the first session she asked me very directly why I was doing the study and why I was interested in posture. Daphne's open and relaxed attitude to the world can be seen in her opinions about the prototypes. She thought both the Buzzy Jacket and Twinkly Lights were helpful because she felt she could improve on both foot position and upper body posture. She did not express a preference for either prototype either in terms of practice focus or feedback type as she could see situations where both might be needed. The smileyometer helped to clarify her opinions a little: she rated the vibrations one point lower in terms of noticing, understanding, effectiveness and overall. Her comments about this focused on not noticing the vibrations and issues when they weren't *"set right".* This may be due to her calibration button never being moved to an easy to use position because after being given the jacket in session 5 she never remember to bring it to any of the following sessions meaning the improvements made to the other jackets were never made to hers.

Her inclination to find fun in everything can be seen in the way that other aspect of the feedback she mainly commented on was how she enjoyed using both types of feedback:

"I enjoyed the lights! It was fun to see them light up although it was bad,"

"[Using the vibrations] it was kind of like playing a game."

Again just as she could find her teacher's criticism funny she also found the lights *"fun"* even though they pointed out mistakes. Her comment about the vibrations being like a game also shows how feedback can be interpreted in different ways.

Enjoying practice was clearly a priority for Daphne whereas as some of the other children were more motivated by working towards achievement goals such as exams or performances. For this reason one might expect her own design to be game-like or overtly rewarding in order to make practice more fun. However her design does not fit this description. In her design to encourage straight bowing she uses visual feedback on a screen. The player's bowing arm would be shown with red lines either side of the bow showing the area which the bow should stay within so the player can see whether they are in the right place. As she developed the idea she also added auditory feedback in the form of a beep to indicate when the bow had strayed outside of the red lines, in case the player had not noticed the visual feedback. The emphasis of this design is on giving information to the player rather than motivating them with ratings or entertainment. However, as Daphne showed in her opinions about the Buzzy Jacket and Twinkly Lights she does not need enjoyment to be explicitly designed into a system she can find entertainment and interest simply in having a new perspective on her playing. This is something that her feedback design would give. Interestingly this design is also very similar to the iMaestro augmented mirror (Ng, Weyde, Larkin, Neubarth, Koerselman, & Ong, 2007).

In another part of the interview she was asked what she would do to improve an aspect of her playing without feedback she talks about thinking about a verbal cue quite similar to the ones her teacher might use to remind her in a lesson:

"Well I'd still focus more on the actual music but then I'd try and keep a bit of my brain saying "fingers, fingers""

Like her teacher's reminders in the lesson it does not contain all the information needed to know exactly what to do with her fingers but is more like a placeholder for something she already physically knows how to do but has difficult remembering (in this case keeping her fingers curved over the string). This kind of reminder is similar to the way the lights worked as feedback whereas the vibrotactile feedback and her bowing design using the screen also give information about what to do to improve and can potentially be used to learn something new.

Daphne's example shows a different motivational perspective on feedback, that of curious inquiry. She found it motivating to be given more information and to be enabled to see things about her playing from a different perspective. While this may not lead to a persistent motivation to practice as the achievement goals of some of the other participants when practice does happen it is intrinsically motivated. When designing for this source of motivation, feedback would need to gradual unfold new views on playing gradually and enable exploration over time in order to keep up interest and curiosity.

Esme

Esme never took her vibrotactile jacket home because her involvement in the study was interrupted by the French exchange, but she had tried the jacket in one of the group sessions. This made it difficult for her to answer questions about the Buzzy Jacket but she gave her first impressions where she could. She felt that the Buzzy Jacket would have been more useful to her than the Footboard and Twinkly Lights because she was more prone to slouching. Her answers to the smileyometer are very similar to those of Daphne who she was interviewed with and it is likely she was influenced a little by her friend's responses although she did spend time considering each response and chose in one case not to answer a question because she didn't feel she had used the jacket enough to judge.

Her smileyometer results suggest that based on her short experience with the vibrotactile feedback she preferred having the lights as feedback. She made very few comments about the two prototypes in the interview. But she did speak about how she would try and improve an aspect of her own playing without feedback which like Daphne involved repeating a verbal cue to help keep the learning focus in mind, but in her case she may say out loud:

"I might keep saying it a lot because that's just what I do."

She felt uncomfortable putting forward a design of her own to help with straight bowing but she did contribute to Daphne's idea of the visualisation on screen with a beep when the player's bow left a certain area.

Appendix F – Ethics and Consent

F.1 Ethics Approval

The studies in this thesis had ethical approval from the Open University (application number 666) and from University College London (application number 3368/001). For the high school study I was CRB checked. All participants were asked to sign a consent form after they had the study explained to them. In the case of the high school study parents were asked for their consent as well – this was mediated by the school.

F.2 Consent Form for MuSense Study (Chapter 5)

I_______________________________________(print name)

agree to participate in this study.

I have had the purposes of the study explained to me.

I have been given the opportunity to ask questions about the study and have had these answered satisfactorily

I have been informed that I may refuse to participate at any point by simply saying so.

I have been assured that my confidentiality will be protected.

I agree that the information that I provide can be used for educational or research purposes, including publication.

I agree that I may be videoed and audio recorded for research purposes.

I understand if I have any concerns or difficulties I can contact:
Rose Johnson r.m.g.johnson@open.ac.uk 07948580029
Janet van der Linden j.vanderlinden@open.ac.uk
Yvonne Rogers y.rogers@open.ac.uk

Additional consents (tick box if you consent)

☐ I agree that still images of me may be used in academic publications and presentations, which may appear online.

☐ I agree that videos of me may be used in presentations, which may appear online.

☐ I agree that still images of me may be used in websites to publicise this research.

☐ I agree that videos of me may be used in websites to publicise this research.

Signed: _________________________________ Date:_____________

If you would like to be kept updated with the progress of this research please give a contact email address:
(this will be kept confidential)__________________________________

F.3 Consent Form Summer School Study

I_______________________________________(print name)

agree to participate in this study.

I have had the purposes of the study explained to me.

I have been given the opportunity to ask questions about the study and have had these answered satisfactorily

I have been informed that I may refuse to participate at any point by simply saying so.

I have been assured that my confidentiality will be protected.

I agree that the information that I provide can be used for educational or research purposes, including publication.

I agree that I may be videoed and audio recorded for research purposes.

I understand if I have any concerns or difficulties I can contact:
Rose Johnson r.m.g.johnson@open.ac.uk 07948580029
Janet van der Linden j.vanderlinden@open.ac.uk
Yvonne Rogers y.rogers@open.ac.uk

Additional consents (tick box if you consent)

☐ I agree that still images of me may be used in academic publications and presentations, which may appear online.

☐ I agree that still images of me may be used in websites to publicise this research.

☐ I agree in principle that video of me may be used in presentations (which may appear online) so long as I am asked consent about the specific video clip which is used.

Signed: ___________________________________ Date:_____________

If you would like to be kept updated with the progress of this research please tick here:

☐

F.4 High School Study (chapter 7) – Student Consent Form

Participant statement

I _______________________________ (participant's name)

agree to take part in this study.

I agree that:

- I have had the purpose of the study and what it involves explained to me.
- I have had the opportunity to ask questions.
- All my questions have been answered to my satisfaction and I know who to contact if I have any further questions about the study.
- I understand that my participation will be video-taped and photographed for research purposes, but that these recording will be kept confidential.
- I understand that I may withdraw from the study without the need to give a reason.
- I understand that by completing and returning this form, I am giving consent that the personal information I provide will only be used for the purposes of this project and not transferred to an organisation outside of UCL and [The High School]. The information will be treated as strictly confidential and handled in accordance with the provisions of the Data Protection Act 1998.

Signed: Dated:

Additional optional consents (tick box and sign if you consent)

[ ] I agree that photographs of myself participating in the study may be used in research publications which may appear online.

[ ]	I agree that short video excerpts of me participating in the study maybe used in research presentations which may appear online.

Signed:	Dated:

F.5 High School Study (chapter 7) – Student Information

Research Study: Investigating New Technologies to Aid Violin Practice – Information for Students

Aims

The purpose of this study is to explore how new sensor technologies can be used to help people practice the violin more effectively. In this study we hope to find out whether the prototypes improve your playing technique and how we can improve the design to make them more successful in the future. To carry out this study we are recruiting violin and viola players from Farlingaye Sinfonia orchestra to meet with us each week and try out the prototypes.

What the Study Involves

The study will run over 3-4 months. If you agree to take part in the study, you will attend an extra rehearsal each week before Sinfonia practice (every Wednesday from 4-5pm). In this rehearsal we will practice the parts you play in Sinfonia in more depth. We plan to video record these practices in order to record you progress.

For the first four weeks the sessions will be just like normal rehearsals so we can find out how you normally play. After four weeks we will bring in some prototypes and use them during the rehearsal. We will also give them to you to practice with at home. The study will then continue for a following 8 weeks or more with us meeting together to rehearse each week.

We are also interested to know how you use the prototypes at home when you are practising. To find this out we'll ask you to keep a practice diary for a week. We will also send you text messages or emails (depending on what you prefer) once or twice a week to find out how things are going.

We'll also give you a video camera for one week and ask you to make a video about using the prototypes. You are welcome to get your parents or friends to help you with this.

We'll also give you a video camera for one week and ask you to make a video about using the prototypes. You are welcome to get your parents or friends to help you with this.

The data we collect about you will be kept secure. It will not be shared with anyone outside of those working on this project. If you don't like any part of this method but you would still like to take part please let us know and we can make special arrangements.

Prototypes

The prototypes in this study use technologies similar to those in the Nintendo Wii to measure how you are playing. Based on these measurements the prototypes will give you feedback about your posture. This feedback will come on while you are playing to remind you to straighten up if you are slouching, or to hold your instrument better. The prototypes give these reminders in different ways. Some show it in lights that attach to the music stand. Others use vibrations similar to a mobile phone to tell you if you are making a mistake. The picture below shows an example prototype which uses lights as feedback.

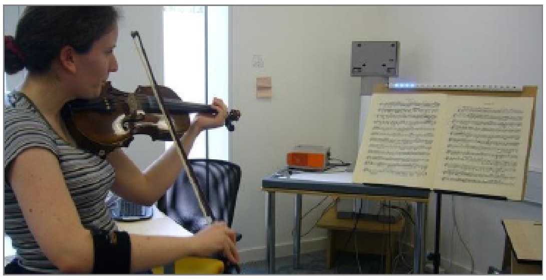

With the technology being so new, there may be a small chance that the prototypes will have the opposite effect to the one we intended. If you feel tense or uncomfortable when you are using the prototypes, please let us know quickly.

Results

At the end of the study we plan to meet with you to show you the things we've found out. We are also planning to write our findings in a research paper to let other researchers know what we have found. You are welcome to have a copy of this.

When we report our findings your data will be made anonymous – no-one should be able to guess who you are from our report. The only case when this is different is if we use photographs and videos in our report because people might recognise you. This is why we have put separate tick boxes on the consent form to allow you to choose whether you are willing for photographs or videos of you to be used when we report our research.

History of the Project

We have done previous research about using technology to help people learn the violin. We did this with a prototype call the MusicJacket which helped people do straight bowing. We experimented with using it in violin lessons with very positive results. If you are interested in finding out more about it, we wrote a research paper about it which you can find here:

http://www.rosejohnson.info/pdfs/vanderLindenCHI2011.pdf

Research Team

Research Team

Rose Johnson

I am a PhD student at University College London and this research will eventually go into my final report. I grew up in [X town] and went to school at [this school], playing violin in the school orchestra. Six years ago I left to study Physics at Swansea University. I completed a Masters there and I am now studying for a PhD in Human Computer Interaction at UCL and The Open University. Last summer I also did a three month internship at Microsoft Research in Cambridge.

I will be the person carrying out the study, if you have any questions about it please contact me. Or if you like you can contact my PhD supervisors Janet and Yvonne.

Email: rose.johnson.11@ucl.ac.uk

Telephone: 07948580029

Dr Janet van der Linden

Janet is a senior lecturer at the Open University and Director of the Pervasive Lab there.

Email: j.vanderlinden@open.ac.uk

Professor Yvonne Rogers

Yvonne is the director of the UCL Interaction Centre.

Email: y.rogers@ucl.ac.uk

F.6 High School Study (chapter 7) – Parental Consent Form

Parent's statement

I _____________________________ (parent's name)

agree that my son\daughter _________________(student's name)

may participate in this study if they consent.

I agree that I have:

Read the information sheet which explains the purpose of the study and what it involves.

I have had the opportunity to ask questions.

All my questions have been answered to my satisfaction and I know who to contact if I have any further questions about the study.

I understand that my child's participation will be video-taped and photographed for research purposes, but that these recording will be kept confidential.

I understand that I may withdraw my child from the study without the need to give a reason.

I understand that by completing and returning this form, I am giving consent that the personal information I or my child provides will only be used for the purposes of this project and not transferred to an organisation outside of UCL and [The High School]. The information will be treated as strictly confidential and handled in accordance with the provisions of the Data Protection Act 1998.

Signed: Dated:

Additional consent options (tick box and sign if you consent):

[ ] I agree that photographs of my child participating in the study may be used in research publications (which may appear online) with their consent.

[ ] I agree that short video excerpts of my child participating in the study maybe used in research presentations (which may appear online) with their consent.

Signed: Dated:

F.7 High School Study (chapter 7) – Parent Information

Research Study: Investigating New Technologies to Aid Violin Practice – Information for Parents

Aims

The purpose of this study is to explore how new sensor technologies can be used to help children practice the violin more effectively. In this study we hope to find out whether the prototypes improve students' playing technique and how to improve the design to make them more successful in the future. To carry out this study we are recruiting violin and viola players from Farlingaye Sinfonia orchestra to meet with us each week and try out the prototypes.

What the Study Involves

The study will run over 3-4 months. Students who agree to take part in the study will attend an extra rehearsal each week before Sinfonia practice (Wednesday 4-5pm). In this rehearsal we will practice the parts they play in Sinfonia in more depth. These practices will be video recorded in order to record students' progress.

For the first four weeks no technology will be used so that we can get a base-line of how the students normally play. After four weeks the prototypes will be introduced and students will be given them to take home. The study will then continue for a following 8 weeks or more.

As well as videoing the rehearsals we are also interested to know how students use the prototypes at home when they are practising. To find this out students will be asked to keep a practice diary for a week. We will also contact them once or twice a week using text messages or email to find out their progress with the prototypes. These particular media have been chosen because they are popular ways for young people to communicate.

Lastly to get a view into home use of the prototypes students will be encouraged to take home a video camera for one week and make a video about using the prototypes. Parents are encouraged to help with this.

The data we collect from your children will be kept secure according to the data protection act of 1998. It will not be shared with anyone outside of those working on this project. If any part of this method makes you uncomfortable but you would still like your son or daughter to take part please do get in touch on the contact details below and we may be able to make special arrangements.

Prototypes

The prototypes in this study use technologies similar to those found in games consoles such as the Nintendo Wii to measure the technique of the student and then give them feedback while they are playing. The idea is that by giving children feedback in real-time they can adapt their movement or posture and avoid practising bad habits.

The main focus for the prototypes will be student's posture. This is something easily overlooked when focussing on playing difficult pieces, but in the long term having good posture can make playing much more comfortable and enjoyable. The aim of the prototypes is to make students more aware of their posture, reminding them to correct it when they begin to slouch. The prototypes give these reminders in various ways. Some show it in lights that attach to the music stand. Others use vibrations similar to a mobile phone to alert the student to a mistake.

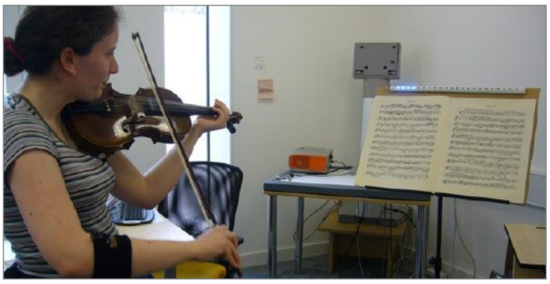

The picture above shows one example of the prototypes we are testing. In this case the sensors are in the sweat band the player wears on her arm. In some cases this is how the sensors may be attached, in others they may be sewn into a jacket or attached to a chair, depending on which elements of posture we focus on.

With the technology being so new, there may be a chance that the prototypes will have the opposite effect to the one we intend. If your child complains of discomfort when playing their instrument or you believe their posture is getting worse rather than better, please let us know quickly so that we can modify our prototypes or halt the study.

Results

At the end of the study we plan to meet with the students to show them some of the things we found. Parents are also welcome to come to this. Additionally we plan to write our findings into a research paper. By doing this we hope to influence future design of musical teaching aids and other areas such as sports practice or physiotherapy. You are welcome to a copy of this. The results reported in the research paper will be anonymous – it should not be possible to identify your child from our descriptions. The one exception is the use of photographs and videos, this is why we have put separate tick box consents on the form to allow you to choose whether you are willing for photographs or videos of your child to be used when we report our research.

History of the Project

This project has been running at the Open University and now at University College London for the last two years. Our earlier prototype the MusicJacket was found to be useful in lesson-time. Now we are keen to find-out whether similar technology may also be useful in practice-time as well. If you are interested in reading our previous paper about the MusicJacket it can be downloaded from the link below, alternatively contact us and we will send you a copy.

http://www.rosejohnson.info/pdfs/vanderLindenCHI2011.pdf

Research Team

Rose Johnson

I am a PhD student at University College London and this research will eventually go on to form part of my thesis. I grew up in [X town] and went to school at [this school], playing violin in the school orchestra. Six years ago I left to study Physics at Swansea University. I completed a Masters there and I am now studying for a PhD in Human Computer Interaction at UCL and The Open University. Last summer I also did a three month internship at Microsoft Research in Cambridge.

I will be the person carrying out the study, if you have any questions about it please do get in contact. Equally my PhD supervisors Janet and Yvonne would be happy to hear from you.

Email: rose.johnson.11@ucl.ac.uk

Telephone: 07948580029

Dr Janet van der Linden

Janet is a senior lecturer at the Open University and Director of the Pervasive Lab there.

Email: j.vanderlinden@open.ac.uk

Professor Yvonne Rogers

Yvonne is the director of the UCL Interaction Centre.

Email: y.rogers@ucl.ac.uk