Adaptive Robotic Carving

Training Methods for the Integration of Material Performances in Timber Manufacturing

Giulio Brugnaro, Sean Hanna

The Bartlett School of Architecture, UCL, 22 Gordon St., London WC1H 0QB, UK

Abstract. The paper presents the developments of a series of methods to train a fabrication system for the integration of material performances in timber manufacturing processes, combining robotic fabrication together with different sensing strategies and machine learning techniques, and their further application within a prototypical design to manufacturing workflow. The training cycle, spanning from the recording of skilled human experts to autonomous robotic explorations, aims to encapsulate different layers of instrumental knowledge into a design interface, giving designers the opportunity to engage with material and tool affordances as process driver. The training methods are evaluated in a series of experiments and design iterations, proving their potential in the development of customized design to manufacturing workflows and integration of material performances, with a specific focus on timber.

Keywords: Material Behaviors, Machine Learning, Instrumental Knowledge, Subtractive Manufacturing.

1 Introduction

The paper presents the developments of a series of methods to train a fabrication system for the integration of material performances in timber manufacturing processes, combining robotic fabrication together with different sensing strategies and machine learning techniques, and their further application within a prototypical design to manufacturing workflow.

Such methods question the linear progression from the design intention to its materialization within current production practices which determines a lack of feedback between the different stages of the process. This forces designers to consider materials as passive receivers of a previously generated ideal form stored in a digital model (DeLanda 2004) and reinforces the separation between the act of designing and making (Carpo 2011). Consequently, design practices can only engage with a limited range of standard manufacturing methods and homogeneous materials. Furthermore, the homogenization of natural material results in heavy industrial processing and material waste.

The goal of the research is to develop a computational framework which allows designers to engage with the properties of heterogeneous materials and the affordances of non-standard tools as process drivers, extending the design moment toward the fabrication stage to explore novel design opportunities. One of the most pressing problems in working such a process is the assessment of inherent variation, represented in this paper by differing wood grain and chisel cuts. There is a range of possible cuts that are feasible, *i.e.* that actually remove material while not damaging either tool or wood, and within this there is a range that may be considered optimal, *e.g.* they remove the maximum amount of material in a given time. Human experts navigate and anticipate this range intuitively, after training and experience. It is a far greater problem for a machine.

This paper investigates how this range and optimum can be quantified and mapped for use in determining appropriate cuts by a robot. The central question is: how accurately are the effects of a given cut predicted based on inputs easily available to the robot? We also investigate whether the domain training inputs makes a significant difference, by comparing data generated by a methodical, parametric, exploration of grain angles and cuts by the computer against data drawn from an expert human user. Is it better for the machine to learn from a skilled teacher, or from its own experience?

2 Context

The methods are based on a series of training and datasets curation procedures, where the instrumental and material knowledge, acquired from both skilled human experts and robotic carving sessions, is captured, transferred, augmented and finally integrated into an interface that makes this knowledge available to the designer.

In this regard, one of the questions is whether is possible to encapsulate, at the least partially, this instrumental knowledge in the technological means for fabrication available to us, making it "*easily accessible, communicable, repeatable, hackable, and transformable*", in the same way, for instance, as 3D modelling software encapsulates knowledge of calculus-based mathematics (Witt 2010). However, in the accumulation of human experience through making and the interaction with materials, there is a tacit dimension that is difficult to capture, formalize and share (Polanyi 1967).

In the history of automation, there have been previous attempts towards this direction, such as the "*Record/Playback*" system to generate machine operations, developed in the 1940s by General Electric. Within this system, a machinist was able to operate a modified machine tool to produce an artefact and get the totality of his motions recorded on a magnetic tape which could be automatically reproduced later by the machine. In those recorded motion was captured not only the gears' mechanical displacement but also the machinist's intelligence, skills and tacit knowledge (Noble 1984, Callicot 2003). More recently, similar methods, focusing on the value of human's action analysis to inform robotic fabrication tasks, have been successfully applied for robotic reproduction of traditional stone surfacing techniques (Steinhagen et al. 2016) and to reconstruct ancient technical gestures associated with the use of tools and the development of the related cognitive functions (Pfleging et al. 2015).

The encapsulation of instrumental knowledge based on sensor measurements of machine operations, rather than human demonstration, finds precedence in a series of industrial manufacturing applications where machine learning techniques, such as Artificial Neural Networks (ANN), have been utilized in a series of machining operation for the optimization of fabrication parameters (e.g. surface roughness, tool wear) and cost reduction (Al-Zubaidi *et al.* 2011). Recently, CITA presented a similar approach to increase the accuracy of robotic incremental sheet forming through the acquisition of scan data and supervised learning (Zwierzycki *et al*. 2017).
The research proposition integrates both approaches for instrumental knowledge encapsulation and combine them in a two-stages robotic training process for the development of custom design to manufacturing workflows, with a specific focus on timber.

3 Robotic Training Methods

This paper tests the accuracy of predictions derived from machine learning on prior example chisel cuts, which are intended in practice to fit into a larger overall workflow (Fig. 1). This consists of a series of training methods which combines the recording of skilled human experts performing subtractive operations with a set of traditional carving tools, such as chisels and gouges, on wooden boards together with autonomous robotic carving sessions.

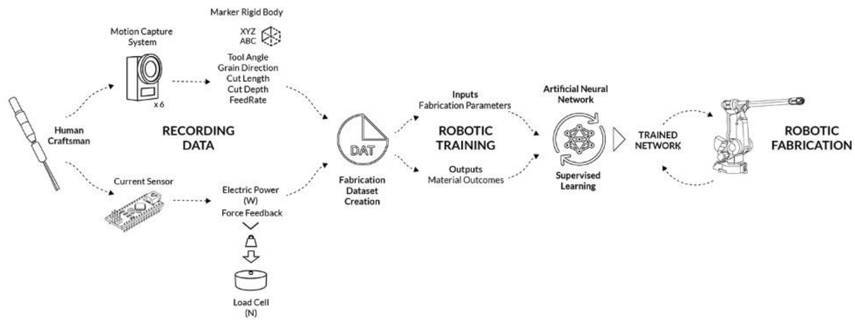

Fig. 1. The robotic training methods are structured around three main stages: 1) Recording fabrication parameters and material outcomes into datasets 2) Encapsulating instrumental knowledge through the training of an Artificial Neural Network 3) Using the trained network to inform a robotic fabrication task.

3.1 Training Tools

For each recording session, the combination of different sensing strategies, such as motion capture cameras and force-feedback sensor, allows collecting fabrication data simultaneously with the performing of the carving operation and compiled it into an ongoing dataset.

Motion-capture cameras are used to track the position and orientation (with a precision of ~0.2 mm) of 3D-printed custom markers applied on the carving tools and workpiece. This allows to stream and reconstruct in real-time in the digital design environment (Rhino3D/Grasshopper) the carving operations and toolpath sequence that generated them (Fig.2).

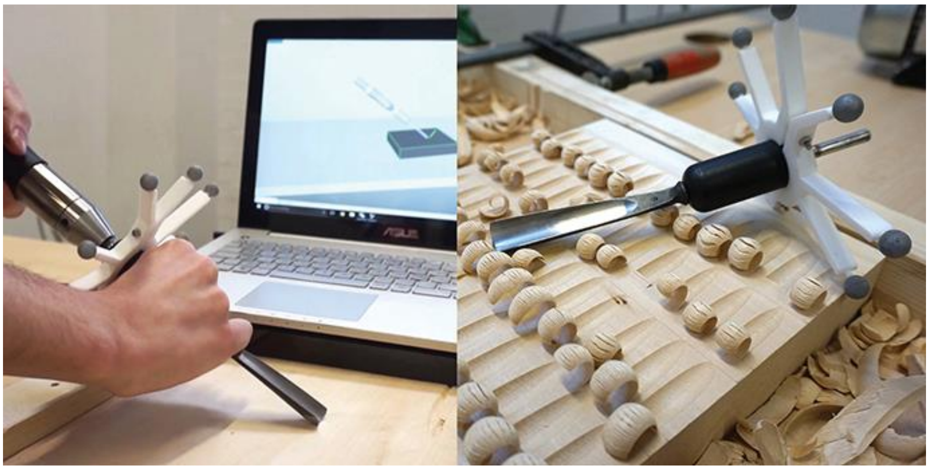

Fig. 2. A system of motion-capture cameras allows to track with high-degree of precision the position and orientation of the fabrication tools and stream this information directly into the design environment.

Following the session, a photogrammetric reconstruction of the training boards is performed to store a precise 3D model of the material outcomes and extract relevant features for the training (Fig.3).

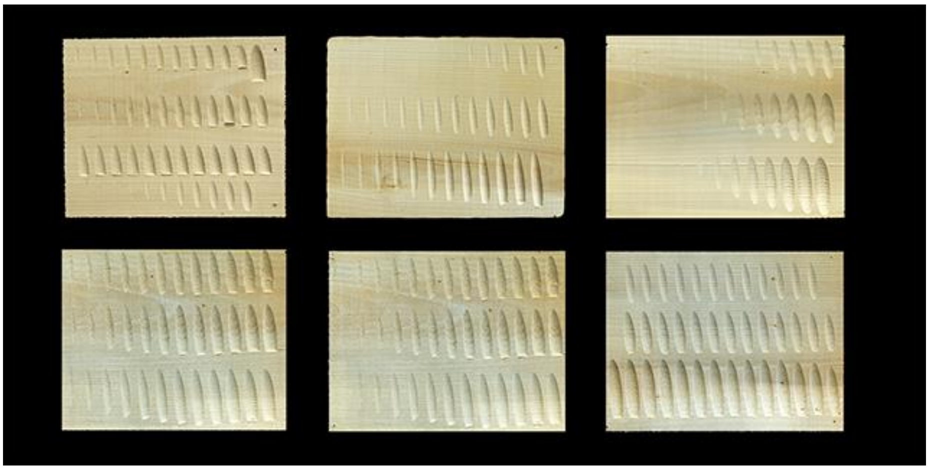

Fig. 3. After the recording session, a photogrammetric reconstruction of each training board is performed to extract relevant features for the training process. In the image, a selection of boards analyzing the outcome of different toolpath carved across the wood grain main direction.

The collected sensor dataset is used within a supervised machine learning procedure, an Artificial Neural Network (ANN) with backpropagation-based learning, whose learning objective is to predict the simulation of a subtractive operation from a user-defined toolpath and a series of fabrication parameters, or conversely, generate a robotic toolpath out of a carved geometry. The training procedure utilizes Tensorflow (GPU version) as machine learning framework combined with the integration of Keras and Scikit-Learn libraries to respectively generate the ANN architecture and measure the performances of the system.

The input parameters are the tool/workpiece angle, tool/grain direction angle, force feedback, feed rate, target cut depth, target cut length. The recorded material outputs are: depth, length and width of the cut. Both the recorded and robotic toolpath are composed of a sequence of target frames, each storing the local parameters and fabrication outcome information and recorded as a single entry of the total dataset.

During the training, the performance of the network is evaluated following a K-fold cross validation (with k=5) procedure, where the dataset is split in k subsets, called folds, and the algorithm is trained on k-1 folds, with each time one of the folds left out to be used to test the system with “unseen” data. The k performance scores obtained at the end of the validation are summarized in the measure of the Mean Absolute Error (MAE) of the prediction and its Standard Deviation (SD).

Starting the robotic training process with the human demonstration allows collecting quickly and efficiently information on how to operate a tool with a specific material, mediated by human experience acquired throughout the years and provide a strong foundation to inform robotic fabrication tasks with a similar set of non-standard fabrication tools and wood species.

The trained network, based on human experts, is used to provide guidance for autonomous robotic training sessions, efficiently narrowing down the search space through the definition of domain boundaries in the selected features rather than arbitrary defining an operational range or relying on a reinforcement learning procedure that would have to learn from a series of potentially dangerous “mistakes” during the fabrication training.

The definition of the search space through human demonstration presents to designers a curatorial approach toward the design of the fabrication process: rather than envisioning a universal machine able to operate any tool on any material, the idea is to tune the system to a very specific set of fabrication affordances and design intentions.

Therefore, the initial dataset is extended through the robotic production of a series of cuts where the parameters investigated are finely interpolated across the training board and the obtained data used for the training of the network informing the actual robotic fabrication task (Fig.4).

Fig. 4. The robotic training sessions allow to perform an in-depth exploration of selected parameters domain through a collection of finely interpolated cuts.

3.2 Design to Fabrication Workflows

The aim of encapsulating instrumental knowledge is its integration into a design interface which makes it accessible to designers and present them the opportunity of using materials behavior into their design workflow as process drivers. Once the network has been trained and the correlations between fabrication parameters and carved geometries are established, it’s possible to translate back and forth between the two sets and customize the network topology to a specific design application (Fig.5).

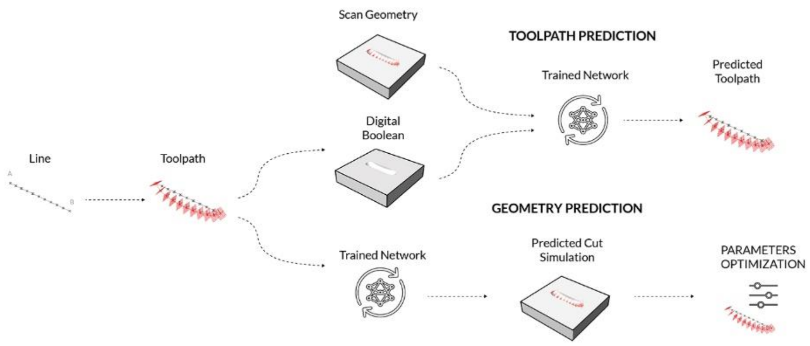

Fig. 5. The instrumental knowledge encapsulated in the trained network is integrated into a design to fabrication workflow based on the opportunity of translating back and forth between robotic toolpaths and carved geometries.

The three main modes of applications explored in the experiments are:

- *From robotic toolpath to simulation of the carved geometry.* While conventional digital Boolean operations result insufficient in calculating the outcome of subtractive operations with non-standard tools on heterogeneous materials, the trained network provides a more accurate simulation based on actual material properties and tool affordances. Designers can directly test how individual fabrication parameters affect the operation result and evaluate how these could be tuned to match their design intention. The prediction could be applied to multiple cuts at the same time, each with different parameters, and used to generate the overall appearance of the cutting pattern.
- *Individual Parameters Optimization.* Utilizing the same set of training inputs and outputs is possible to create labels (or Boolean flags) to predict a series of event thresholds based on sets of fabrication parameters, such as the successful removal of material or the correct extraction of the tool from the workpiece. Moreover, additional labels could be created by the designer to describe formal preferences (e.g. surface roughness, edges definition), curating the training dataset along a specific design direction. The event threshold is predicted using a ANN for binary classification and the prediction accuracy is evaluated with the same cross validation method previously described.
- *From carved geometry to robotic toolpath*. Extracting fabrication data out of the scanned model of a previously carved workpiece to reconstruct the robotic toolpath that has generated it. Alternatively, the same method could be applied starting from a digital geometry obtained through a subtractive Boolean operation as a way of matching a formal design intention in the fabrication stage.

4 Experiments and Results Discussion

As part of their development, the training methods have been iteratively tested to inform the robotic fabrication of a series of carved panels using different wood species and carving tools (Fig. 6).

Fig. 6. In the robotic fabrication stage, the trained network is used to inform the robotic toolpath to carve the previously simulated geometries.

One of the key things evaluated has been how the integration of human instrumental knowledge compares to the robotic training process. To perform this evaluation, two different training cycles have been set up, one starting with the recording of skilled human expert manually performing carving operations on a series of boards, the other directly with the industrial robotic arm (ABB IRB1600) generating a collection of cuts within an arbitrarily defined range of fabrication parameters.

In the first experiment, extending the findings of previously published work by the authors (Brugnaro and Hanna, 2017), the craftsman created an initial collection of 50 carving operation on a series of European Lime wood boards using a traditional carving gouge (Stubai 9/20). The dataset has been generated concurrently with the carving itself following the methods previously described. The human expert, thanks to his experience, has been able to avoid operations where the tool was not actually cutting the material or, on the opposite, digging too much into it and consequently forcing to stop the operation. Moreover, once individuated a preferred carved result, the cutting goal has been adjusted through human judgement to achieve a similar qualitative result in different lengths, depths and wood grain directions.

In the second stage of the experiment, the initial human-based dataset has been used to narrow-down the parameters domain and systematically explore it with the industrial robot through a collection of 215 cuts.

While the final prediction goal of the training cycle is the simulation of a carved geometry according to a set of given fabrication parameters, the first stage, based on human expert data, is aimed toward the generation of robotic toolpaths for additional subtractive operations that would extend the initially acquired dataset.

In the specific, the dataset of the first stage has been used to: 1) Analyze how the domain boundaries of the key parameters of tool/surface angle variation and “depth” profile of the toolpath changes according to different cut lengths, widths and wood grain directions. 2) Train a network and use it to predict such fabrication parameters to generate the robotic toolpath for the 215 operations within the dataset defined boundaries. The prediction error for the toolpath generation stage was the following: Depth: MAE = 0.56 mm, SD = 0.12mm; Tool Angle: MAE = 2.12°, SD = 0.35°.

In the second experiment, counting the same number of cuts and using the same type of wood and carving gouge to compare it to the previous, the range of the fabrication parameters to explore has been defined without the guidance of a human expert interacting with the materials and tools but directly with the robotic operations stage. The focus of the training session has been on the variation of the material outcomes in respect of the angle between tool, workpiece surface and grain direction applied to different length of the cuts.

It’s important to note that in the second experiment, the user applied his undirect intuition and understanding of the task in the programming of the robotic actuation, while for the first experiment direct real-world fabrication data have been used to inform it.

The output of the experiments are two networks trained with datasets counting the same number of cuts but different prediction boundaries. As a consequence, they perform differently in the prediction of the material outputs (length, depth and width relative to each individual target frame) necessary to simulate the fabrication outcome of the chosen operation.

In the specific, the two-stages (human+robot) network has been trained toward a specific design intention and performs better within its narrowed-down fabrication parameters domain (Depth: MAE = 0.35 mm, SD = 0.04 mm; Width: MAE = 0.83 mm, SD = 0.13; Length: MAE = 0.28mm, SD = 0.05mm).

The single stage experiment (robot only), while in general performing worse than the former (Depth: MAE = 0.53 mm, SD = 0.12mm; Width: MAE = 1.02 mm, SD = 0.32; Length: MAE = 0.78mm, SD = 0.32) and resulting less efficient for some operations described below, is able to cover a more extended and generic range of parameters prediction and could result more useful in those cases when the design task is not strictly defined from the training stage.

The dataset generated in the second experiments presents some set of parameters which does not generate any material removal on the board or damage the tool not allowing to complete the operation. In this case, an important step has been introducing two different labels for the definition of event thresholds: 1) “Tool Damaging” and 2) “Material Removal”. Once these labels have been assigned, it’s been possible to train the network for a binary classification process of the two event states (Fig. 7). The accuracy of the prediction for the two labels were respectively of 92.45% (Standard Deviation: 4.96%) and 84.58% (Standard Deviation: 3.20%). Combining the binary classification prediction with the trained network of the second experiment has been possible to efficiently tune the individual parameters and filter out dangerous or inefficient robotic operations (Fig. 8). The prediction of the range of successful cuts allows to explore its boundaries with confidence and optimize individual parameters to achieve a specific material outcome or increase the efficiency of the process, maximizing, for instance, the removal volume (Fig. 9).

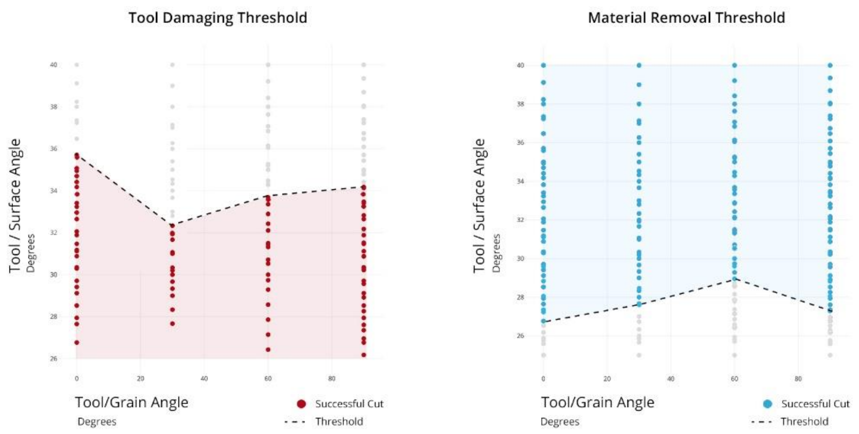

Fig. 7. "Tool Damaging" and "Material Removal" event thresholds are predicted with an ANN for binary classification and allow to optimize individual fabrication parameters to avoid potentially dangerous or inefficient operations.

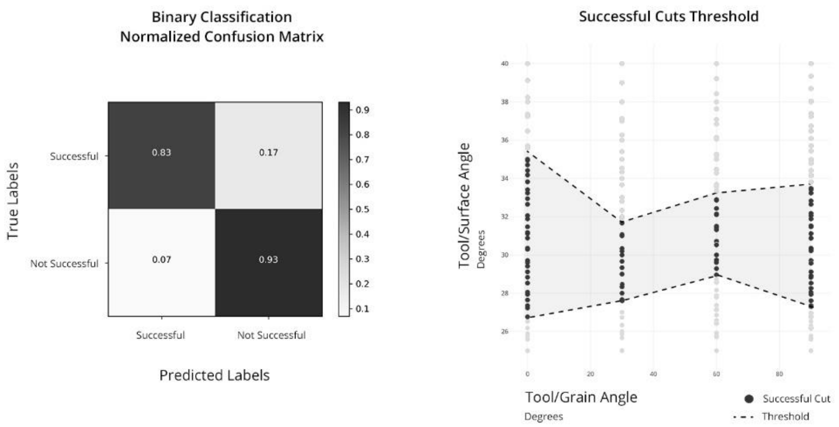

Fig. 8. The event threshold for "successful" cuts is predicted combining the previous event labels, defining the range of optimal cuts. A normalized confusion matrix shows the prediction rate for the "Successful" and "Not Successful" labels, with an accuracy of 83% and 93% respectively.

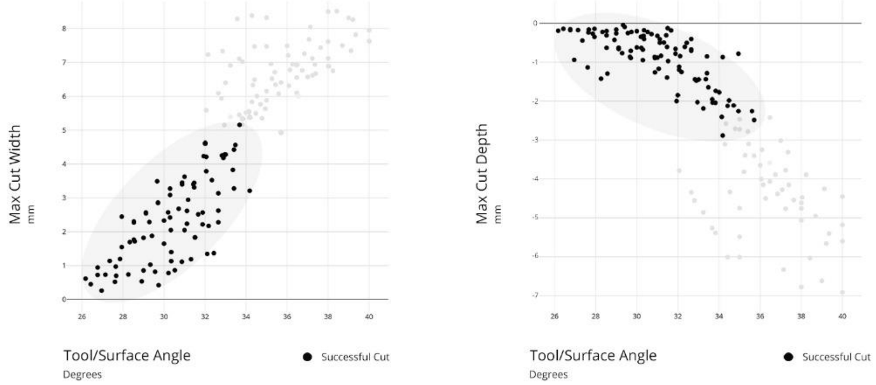

Fig. 9. The tool/surface angle value at the beginning of the cut affects significantly the successful result of the operation. Such parameter could be optimized toward a specific material outcome, such as maximizing the depth and width of the cut to increase the efficiency of the subtractive process.

Finally, the dataset generated during the robotic training have been used to reverse the network topology of the previous experiments and predict the tool/workpiece angle parameter and input cutting length necessary to reconstruct a toolpath used to generate a carved geometry. The prediction error was respectively 1.69° (SD = 0.42°) and 3.37 mm (SD = 1.23 mm) and the overall process performs slightly worse in the reconstruction of the toolpath than the network trained in the opposite direction for geometry simulation, probably due to hardware noise and imprecisions in the physical measurements, confirming what already pointed out by Zwierzycki *et al.* (2016) in their research.

5 Conclusions

Focusing on robotic subtractive fabrication with timber as the main case study, the paper presents the potential of machine-learning strategies for design to manufacturing applications, as a way to explore novel design opportunities through the integration of actual material properties and tool affordances for simulation purposes (Fig. 10, Fig. 11).

The prediction rates of the trained networks presented in the experiments discussion suggest that it's possible to accurately simulate the results of subtractive operations based on a given set of fabrication parameters and use the encapsulation of such instrumental knowledge to translate back and forth between robotic toolpath data and geometry prediction.

The experiments focused on two different workflows for the generation of fabrication datasets and curation of the training process, examining whether is better for the robotic system to learn from a skilled human expert or autonomous training sessions. The lower prediction error of the network based on the combination of human and robotic training should be attributed to the steering action of the human expert, operating within a narrowed down parameters range, which excluded inefficient or dangerous cuts, and leading the robotic training sessions to gather data only within an optimal fabrication range. The more systematic and wider range of the second experiment doesn't allow to have the same resolution in the more relevant areas of "successful" cuts, resulting in lower prediction performances.

To conclude, the manipulation of knowledge across distinctly operating domains such as human making and industrial robotic manufacturing presented the opportunity to develop an approach for human-machine interaction which questions current industrial method of knowledge transfer and will be further explored in future steps of the research.

Fig. 10. The training process allows to tune the fabrication system to a specific set of tools and type of wood to explore design opportunities that would become evident only in the fabrication stage.

Fig. 11. Combining together prediction of cuts with different fabrication parameters, it's possible to simulate the final carved geometry in several pattern configurations.

References

1. Al-Zubaidi, Salah, Jaharah A. Ghani, and Che Hassan Che Haron. 2011. "Application of ANN in Milling Process: A Review." Modelling and Simulation in Engineering 2011: 9
2. Brugnaro, Giulio, Hanna, Sean. 2017. "Adaptive Robotic Training Methods for Subtractive Manufacturing". In ACADIA 2017: Disciplines and Disruption, Proceedings of the 37th Annual Conference of the Association for Computer Aided Design in Architecture (ACADIA). Cambridge (MA): 164-169.
3. Callicott, Nick. 2003. "The Tacit Component and The Numerical Model: Representation In Computer-Aided Manufacture And Architecture". *The Journal Of Architecture* 8 (2): 191-202.
4. Carpo, Mario. 2011. *The alphabet and the algorithm*. Cambridge, Mass: MIT Press.
5. DeLanda, Manuel. 2004. "Material Complexity." In Digital Tectonics: 14–21.
6. Noble, David F. 1984. *Forces Of Production*. New York: Knopf.
7. Pfleging, Johannes, Marius Stucheli, Radu Iovita, and Jonas Buchli. 2015. "Dynamic Monitoring Reveals Motor Task Characteristics in Prehistoric Technical Gestures." In *PLoS ONE 10 (8)*: 1–20.
8. Polanyi, Michael. 1967. *The Tacit Dimension*. Garden City, N.Y.: Anchor Books.
9. Steinhagen, Gregor, Braumann, Johannes, Brüninghaus, Jan, Neuhaus, Matthias, Brell-Çokcan, Sigrid, Kuhlenkötter, Bernd. 2016. "Path planning for Robotic Artistic Stone Surface Production". In *Robotic Fabrication in Architecture, Art and Design 2016*. Springer: 122-135.
10. Witt, Andrew J. 2010. "A Machine Epistemology in Architecture. Encapsulated Knowledge and the Instrumentation of Design." In *Candide. Journal for Architectural Knowledge 3 (3)*: 37–88
11. Zwierzycki, Mateusz, Paul Nicholas, and Mette Ramsgaard Thomsen. 2017. "Localised And Learnt Applications Of Machine Learning For Robotic Incremental Sheet Forming". *Humanizing Digital Reality*, 373-382. doi:10.1007/978-981-10-6611-5_32.