Integrating Ontologies and Argumentation for decision-making in breast cancer

Matthew Hardy Williams

A dissertation submitted in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

of the

University College London.

Department of Computer Science

University College London

2008

I, Matthew Hardy Williams, 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

This thesis describes some of the problems in providing care for patients with breast cancer. These are then used to motivate the development of an extension to an existing theory of argumentation, which I call the Ontology-based Argumentation Formalism (OAF). The work is assessed in both theoretical and empirical ways.

From a clinical perspective, there is a problem with the provision of care. Numerous reports have noted the failure to provide uniformly high quality care, as well as the number of deaths caused by medical care. The medical profession has responded in various ways, but one of these has been the development of Decision Support Systems (DSS). The evidence for the effectiveness of such systems is mixed, and the technical basis of such systems remains open to debate. However, one basis that has been used is argumentation.

An important aspect of clinical practice is the use of the evidence from clinical trials, but these trials are based on the results in defined groups of patients. Thus when we use the results of clinical trials to reason about treatments, there are two forms of information we are interested in - the evidence from trials and the relationships between groups of patients and treatments. The relational information can be captured in an ontology about the groups of patients and treatments, and the information from the trials captured as a set of defeasible rules.

OAF is an extension of an existing argumentation system, and provides the basis for an argumentation-based Knowledge Representation system which could serve as the basis for future DSS. In OAF, the ontology provides a repository of facts, both asserted and inferred on the basis of formulae in the ontology, as well as defining the language of the defeasible rules. The defeasible rules are used in a process of defeasible reasoning, where monotonic consistent chains of reasoning are used to draw plausible conclusions. This defeasible reasoning is used to generate arguments and counter-arguments. Conflict between arguments is defined in terms of inconsistent formulae in the ontology, and by using existing proposals for ontology languages we are able to make use of existing proposals and technologies for ontological reasoning.

There are three substantial areas of novel work: I develop an extension to an existing argumentation formalism, and prove some simple properties of the formalism. I also provide a novel formalism of the practical syllogism and related hypothetical reasoning, and compare my approach to two other proposals in the literature. I conclude with a substantial case study based on a breast cancer guideline, and in order to do so I describe a methodology for comparing formal and informal arguments, and use the results of this to discuss the strengths and weaknesses of OAF. In order to develop the case study, I provide a prototype implementation. The prototype uses a novel incremental algorithm to construct arguments and I give soundness, completeness and time-complexity results. The final chapter of the thesis discusses some general lessons from the development of OAF and gives ideas for future work.

Acknowledgements

I would like to start by thanking Cancer Research UK for their generous funding of the work in this thesis, as well as associated work over the time. In addition, I would like to thank many of the staff of Cancer Research UK, particularly Yvonne Harman, Ava Yeo and Nancy Hogg for all their support, as well as the members of my thesis committee, Richard Begent, Peter Parker and Ian Tomlinson, for their feedback on the early stages of this work.

The substantial part of this work was carried out while based at the Advanced Computation Lab at Cancer Research UK, and I must thank John Fox for originally inviting me to the lab, as well for many discussions over the years, and feedback and criticism on the work in progress. I must thank many people in the lab for conversations and ideas over the years, but particular thanks go to Sanjay Mogdil, Scott Hartop and Rory Steele, as well as Liz Black and Ali Rahmanzadeh, with whom I shared an office for so long, and were always willing to discuss even the more ridiculous ideas.

The rest of the work was done in the Computer Science Dept. at UCL, and my thanks to Patricia Fenoy, JJ Giwa and Naomi Jones for all their help with the practical matter of studying for a PhD. The majority of the ideas in the thesis were developed in discussion with and under the supervision of Tony Hunter at UCL, and I am grateful for his repeated reading and criticism of the work over the years, as well as taking a chance on supervising me in the first place. My thanks also to Nikos Gkorogiannis, whose advice and suggestions on the work in the thesis and related work has been consistently stimulating and interesting.

My final thanks must go to my family, especially my wife without whose support none of this would have been possible, and my children, both born during the course of the thesis, without whom I would have had a great deal more sleep.

Contents

Chapter 1

Introduction

As living and moving beings, we are forced to act .... [even when] our existing knowledge does not provide a sufficient basis for a calculated mathematical expectation

JM Keynes, 1936 [47]

Medicine is a humanly impossible task

A Rector, 2001 [65]

This chapter provides the background to both the clinical and technical aspects of the thesis. I start with an introduction to cancer biology, and breast cancer in particular, and then discuss the problem of providing high quality medical care. As I show, part of this problem is related to the provision of information, and one solution has been the use of Decision Support Systems (DSS). I introduce one formal basis for such systems - argumentation. I give a brief review of some different approaches to argumentation, and show that while argumentation has some advantages as the basis for DSS, current approaches have some weaknesses which can be resolved by considering the requirements for a Knowledge Representation formalism. I use these weaknesses to motivate the work in this thesis, and conclude with a set of requirements for a new argumentation-based knowledge representation formalism.

1.1 The Clinical Problem

Ms. Jones is a 55 yr old woman with left-sided adenocarcinoma of the breast. The tumour has a maximum diameter of 18 mm, and is found to be grade 2 and ER/PR positive, but HER-2 negative. She had a wide local excision of the tumour and 10 axillary lymph nodes removed, one of which contained metastatic disease.

What course of treatment would you advise her?

As a doctor, the patient described above is typical of many patients I have seen in clinic, and in some ways represents many of the challenges facing us in clinical medicine, and oncology in particular. Over the last 30 years, we have developed an increasingly rich understanding of the biology of disease, and the impact of different treatments; unfortunately, as the second quote suggests, this has not made the clinician's job easier. We face a variety of problems - there are an increasing number of possible treatments, an increasing amount of information about each treatment, and an increasing understanding of the nature of the disease (which tends to make things more, not less, complex). At the same time, we face an (increasing) number of patients, who need treating now, even in the light of incomplete and inconsistent information. To decide which treatment options we would recommend in this case, we need to consider not just what treatment the patient has had (surgery), but also the exact nature of the tumour, the likely benefits of different treatments and the importance that she places on both the benefits and side-effects of the treatments.

1.2 Clinical Background

Cancer is the clinical manifestation of the uncontrolled growth and replication of cells. This occurs via a variety of mechanisms (e.g. exposure to radiation, exposure to chemicals, incorporation of viral DNA) all of which alter cellular DNA. As a result, cells lose control of the cell-cycle. Most often, this leads to cell death, but occasionally the cell begins to multiply, unhindered by normal restrictions. All cancers have their basis in normal tissues, and the site and type of this tissue is important in determining the type of cancer that will develop, and hence the type of treatment that is likely to be effective. In addition to classification by tissue of origin, however, we can also classify tumours based on how abnormal their cells appear under a light microcope (their ‘grade’ on a scale 1-3, where ‘1’ indicates abnormal but still with many normal features and ‘3’ represents grossly abnormal), what type of cells they are (their histology), the extent of their spread to either lymph nodes or surrounding tissue and the size of the primary tumour (their stage). In addition to local spread, cancers also have a tendency to spread to distant organs (metastasis). The presence of metastasis is an important prognostic factor, but while we may be able to detect gross metastatic disease, the presence of very small groups of cells is impossible to detect. This therefore often leads us to having to treat some patients on the presumption that they are likely to have metastatic disease, even if we cannot find clear evidence of it.

In addition to mutations in the genes that control the cell-cycle, cancerous cells often have mutations in other genes. These mutations may result in the production of certain proteins expressed on the cell surface, and these proteins can be detected and their level quantified. These proteins may then become the target for subsequent therapies - for example, Tamoxifen (and related drugs) target the Oestrogen receptors that are found on many breast-tumours, while Herceptin targets the HER-2 receptor. In this case, however, the presence of both receptors plays a dual role: Not only do they determine the efficacy of the drugs [74, 22], they also give prognostic information. Thus, patients with breast cancers that are strongly HER-2 positive have a worse prognosis that those with tumours that are HER-2 negative, but will respond to drugs (such as Herceptin) that target HER-2 receptors.

When we descibe a tumour, therefore, we report a whole variety of information: Its location, its histological type, its size and spread (so called TNM stage), as well as the presence of certain cell-surface proteins. Which information is relevant to which tumour depends on many factors. These same factors are widely used in the inclusion and exclusion criteria for clinical trials and form an important guide as to what knowledge we consider important.

1.2.1 Breast Cancer

Breast Cancer is one of the most common cancers in the Western World. It is the most common non-skin cancer in women in the UK and USA, and accounts for approximately 1/3rd of cancers in women, with the risk of developing breast cancer at some point in a woman's life being approximately 1 in 10. Some 36 000 cases are diagnosed each year in the UK, of whom about a 1/3rd will die from the disease [54]. Consequently, there has been a considerable amount of research focused on breast cancer, and death rates have fallen over the last 10 years[63].

Treatments for Breast Cancer: The mainstay of treatment for breast cancer remains surgery and radiotherapy, [67] with hormone and chemotherapy treatment often used for presumed metastatic disease. One of the advantages of surgery is that, as well as removing any local disease, a sample can be taken from the axillary lymph nodes. These are a common site of lymphatic spread for the cancer, and their removal not only removes any spread that may have occurred, but also allows analysis of the nodes to describe the degree of spread. The two main aims of treatment are to provide local control of, and to prevent premature death from, disease. However, to achieve this aim (especially in high-risk patients) requires the use of complex and toxic chemotherapy regimes, often augumented by the use of other drugs. These treatments (like all treatments) have an associated morbidity and mortality, and so we aim to select those patients who will benefit most from such treatments. In general those at high risk are treated aggressively, while those at low risk are treated less aggressively. This restricts the side effects of intensive treatment to those patients who would benefit most and allows more efficient use of resources. This choice of therapy for post-surgical breast cancer patients is made by a multi-disciplinary team of surgeons, oncologists, histopathologists, specialist breast care nurses and others in a joint meeting (the MDM, or multi-disciplinary meeting), and it is this arena that will provide the motivation for my work.

1.2.2 The Problem of Care

To illustrate the scale of the information problem, according to the Index Medicus, more than 500 clinical trials in humans on breast cancer were published in 2004, and some of those 500 papers will already be contributing to decisions about which treatments patients should be offered. However, the number of these publications has lead to a set of new problems. Given the rate of change in medical knowledge, how can clinicians stay up to date? How can we ensure that current practice is in line with best practice? How can we ensure that local variations in care do not affect patients adversely? Over the last 15-20 years, there have been a set of responses from the medical profession to try and address these problems (see 1.3 for details). Unfortunately, this is still not enough, and two recent reports have highlighted the differences in care between settings. The Cancer 2025 report [18] suggested than if we were able to close the gap in care between the best and worst centres in the UK, we would reduce mortality by 10%, a figure that exceeds any recent single advance in treatment. These problems are not restricted to the UK and similar problems exist in the USA.

It is not the only problem in knowing what care to deliver - a recent US report from the Institute of Medicine [49] suggested that adverse events in medical care caused 44 - 98 000 deaths per year in the USA, which is a problem with knowing *how* to deliver care safely. However, this thesis will focus on the first problem, which is that we do not know *what* to do, and to solve it requires us to understand and manage medical knowledge.

1.2.2.1 The Problem of Knowledge

To be able to manage medical knowledge requires that we know what we are trying to understand and represent, and we must be clear on what this is. One commonly used distinction is between data, information and knowledge [1]. In this classification, data is the lowest level of the hierarchy, such as the string representing a person's name, or an integer representing their age. Information is this data in context, such as knowing that some string refers to a person's name, not their address. Knowledge might be knowing that if we are to write a formal letter, we should use their surname, and their first name if informal. In this division, we are interested in knowledge (accepting that we have to use data and information to represent the knowledge), but we can divide the types of knowledge we have further. For example, we can distinguish the knowledge that aspirin is a non-steroidal inflammatory drug (NSAID) from the knowledge that NSAIDs are useful in treating pain, and also affect blood clotting. In turn, that should be distinguished from the knowledge that aspirin is useful in treating patients with heart attacks, and is part of the protocol in one particular hospital for the treatment of heart attacks. This is of course different from the knowledge that patients with heart attacks need admitting to the Coronary Care Unit.

My purpose here is not to try to elucidate all the different types of medical knowledge, but just to illustrate that it takes many different forms. The development of a system that is able to handle all the different sorts of medical knowledge is outside the scope of my thesis. Instead, I shall concentrate on two forms of medical knowledge. I referred to the number of clinical trials above, and so I focus on developing a system to represent and reason with the results of clinical trials. However, to do this, I will make use of knowledge about how the terms used to describe the patients, interventions and results of trials relate to each other. Discussion of the results of clinical trials are part of the Multi-disciplinary Meeting (MDM) and are seen as important, but from a medical perspective, the information about the relationships between types of treatments is largely seen as uncontroversial background knowledge. This distinction will be important for the work in the thesis.

On the one hand, we have knowledge about terms in the clinical domain - drugs, patient groups, outcomes. These are explicitly described, often quite static concepts (the nature of aspirin is defined in relation to its chemical structure; the nature of the group of women with early breast cancer is carefully defined in terms of stage of disease). As such, we can describe the relationship between each group relatively easily, and such relationships are not interesting to those who work in the area - they are a given of their work. In contrast, the knowledge generated from clinical trials is far less well structured, is often in conflict, may be only partially applicable and may often increase uncertainty - as the number of drugs grows, so does the number of possible interventions. There are attempts to resolve some of these conflicts (e.g. in the publication of reviews and meta-analyses) but these are dependent on there being enough evidence from different sources to resolve, and so lag behind the production of the knowledge. In addition, increasing understanding of the basic science behind disease and treatment also allows us to generate speculative, but often plausible inferences even before clinical trials are done. The interpretation of the results of many trials depends on the values that are attached to certain outcomes. Funders may prefer lower-cost treatments, while patients may prefer more effective treatments, or ones they can take at home. Therefore the question of which is the 'best' treatment is difficult for at least two reasons, involving the interpretation of information that is both conflicting and value-dependent.

At the same time, clinicians are faced with the same over-riding question: What is to be done, today, for this patient? Such a question negates the common response to uncertainty ('Elict more information') and demands that we commit to some course of action. One option might be to simply do what has been done before, until the effect of new treatments has been resolved; another might be to pick a new treatment at random; another might be to allow the patient to decide. All of these are, in some cases, defensible, and in others less so; while we might be happy with exisiting treatments where it is generally effective, if the existing treatment is poor, then we may tend to novelty; random choices of treatments satisfies our novelty, but removes even the pretence of rationality (although admits the degree of equipose necessary for clinical trials); allowing patients to decide assumes that they are able to both exercise their judgement effectively, and are also able to dispassionately weigh their own interests against those of the wider community.

1.3 Approaches to the problem of knowledge

The reaction by the medical profession has included the development of practices such as a requirement for Continuing Medical Education, Evidence-Based Medicine (EBM), the publication of structured abstracts, and a gradual move towards sub-specialisation (referred to as site-specialisation in oncology). These aim to ensure that clincians regularly attend educational meetings, improve the standard of the information on which they base their decisions, to present that information in a more digestible form and to reduce the amount of information one clinician is expected to handle.

1.3.1 Evidence

Evidence-Based Medicine has become a major intellectual and practical force in clinical medicine over the last twenty years. We can distinguish two main strands in EBM: The first relates to the principle of examining why we do what we do [68]; the second relates to how we handle that information. The only approach that has become widely accepted amongst clinicians in the second area is Bayes' Theorem for diagnosis and expected utility theory for decision making. This has had a substantial impact, especially in the area of diagnosis and test choice, where the ability to quantify pre-test probability and test performance has allowed robust reasoning about which patients we should and should not subject to certain tests, and the merit of one test over another. However, many clinicians have problems with the strictly Bayesian approach to EBM, and this has generated a great deal of literature, some of which has identified some very pertinent weaknesses in the use of Bayes' theorem in a clinical setting. There are particular weaknesses in the application of data that may be widely accepted as 'true' but which lacks a sound statistical basis - and it is often this sort of information which clinicians value. Having said this, it remains the only widely acceptable model in use by clinicians and clinical academics.

Outside of the EBM approach, there has been a steady interest in the nature of 'Evidence', especially in the legal profession. Early work by authors such as Wigmore [77] on the diagrammatic representation of legal evidence has been supplemented by more recent work by authors such as Schum [69] and Dawid [25], both of whom use a Bayesian approach to evidence. Under this approach, an inferential link is established as a set of probabilistic relationships, such that some piece of evidence (say, the presence of a blood stain) is taken as being suggestive of the previous presence of a person whose DNA matches the blood, which in turn is evidence of the fact that they had the opportunity to commit the crime, and so on. Although this approach is interesting, especially in the common patterns of evidence identified by Schum, it only allows the integration of conflicting evidence as represented by a change in the probability function of some variable. The problem is that when dealing with mutually exclusive information, it makes little sense to attempt to 'integrate' both into the result of some outcome variable; instead, either one is correct, or the other is, and it matters which one you believe as to where you will estimate the probability to lie. This can be handled by a Bayesian approach if one makes the evidence conditional upon it being believed (i.e. the blood-stain is evidence if the presence of the blood-stain is believed), but this is an immensely cumbersome approach. There may be a multitude of reasons for disputing the link between the bloodstain and the presence; to enumerate these in the conditions of the evidence is very time-consuming and difficult.

1.3.2 Decision Support Systems

Despite the variations in care, there is an internal tension between the drive towards increasing specialisation and the demands and needs of patients to be treated nearer home, at less cost, as outpatients, and in a more flexible manner. This internal tension means that the 'obvious' solution of concentrating care in a few high volume centres has limits. Instead of this, we need to focus on extending the number of centres that are able to offer such high quality care, and the range of staff able to offer such care. In addition, there is a recognition that translation of evidence into practice remains a substantial problem. Current approaches include the use of many different techniques, including focused teaching sessions, the development and dissemination of explicit guidelines, the use of care pathways and, increasingly, the development of Decision Support Tools (DSTs). DSTs encompass a range of technologies from paper-based flowcharts to semi-automatous computer programs and I will reserve the term Decision Support System (DSS) for "active knowledge systems which use two or more items of patient data to generate case-specific advice" [79]. The first suggestion that logic could provide a basis for such systems is generally credited to McCarthy [53], although the paper makes no mention of specifically medical uses. In clinical medicine the first substantial development was my Shortliffe in Stanford [24], and the first use was in Leeds by de Dombal [26]. Despite much initial enthusiasm, none of these techniques has generated the hoped-for advances in care, and the results of the DSS, whilst appearing impressive, have often failed to have a substantial clinical impact. The reasons for this are many and varied, and encompass a variety of political, social, technical and scientific reasons, but despite these disappointments some recent meta-analyses of DSS impact have suggested that they can have a positive impact on several aspects of patient care (although their effect may be more in promoting adherence to guidelines rather than in improving outcomes) [45, 41]. Other studies have shown that even widely deployed systems can have little effect [29]. There are also substantial issues around usability and user-acceptability. This does not invalidate the idea of DSS - instead, it demands a better understanding of the failures, and ways to overcome them. Given these two problems - the need for better ways of using the knowledge we have, and the relative failure of existing techniques - DSS continue to appear attractive. However given the failure of existing DSS, new developments in inference and decision-making, such as argumentation, may be of value. I give a more comprehensive introduction to argumentation in the next chapter, but for now we can regard argumentation as being a technique for reasoning about and with potentially conflicting information by constructing lines of argument that support different, possibly conflicting claims.

1.3.3 Argumentation-based DSS

There are several possible formal bases for DSS, and in addition to the evidence on the impact of DSS in general, there is some specific evidence that argumentation-based DSS tools are effective [44, 76, 59], and this is the strand of work I shall pursue. However, all of these implemented argument-based DSS are based on the Logic of Argumentation (LA) work of Fox, Sutton, et al [50]. This work was some of the earliest in the field, and when viewed in the light of later work some clear weaknesses are apparent. The main problem is that arguments are directly authored, rather than being constructed from smaller units of knowledge, and conflicting arguments are resolved on the basis of the number of arguments for and against a point, rather than any other idea of argument interaction being considered. In addition, there is no committment to any strong data model, and certainly no idea of using some externally defined data model, which makes integrating the system with external data sources (such as an electronic patient record) hard. Because of this, there is a tendency to author bespoke knowledge bases for each use of the system, and there is fundamentally a mismatch between the way that knowledge is presented in the domain (in the form of clinical trials) and the way it is encoded in the system (as arguments). The end result of this is that there are some major problems with scalability and maintainability of the knowledge bases used in such systems, which are also difficult to integrate with existing medical data sources.

1.3.4 Towards a new argument-based approach

So far, I have given some background to cancer, and breast cancer in particular, and introduced the idea that it is difficult to deliver high quality care. Part of this difficulty is a problem of knowledge, and I have reviewed some of the existing approaches to the problem. One of these approaches is the use of DSS, and argument-based DSS represent a small but interesting group of implemented DSS. However, there are some weaknesses in the underlying formalism used in such systems, and these weaknesses provide the basis for this thesis. I shall concentrate on developing an argumentation formalism to represent and reason with clinical knowledge. Specifically, I am interested in representing and reasoning with the results of clinical trials, although there is some associated background knowledge that we may need to make use of as well. This background knowledge should be the same for our formalism as other approaches, however, and so provides an obvious point where we should make use of external data sources and knowledge.

1.3.5 Knowledge Representation

There have been man different approaches to argumentation over the last 15 years, but the emphasis so far has been on the logical basis of argumentation. Across all of the work there has been a concentration on *how* to argue - that is, what constitutes an argument and how we decide if one argument defeats another, rather than *what* it is that they are arguing about. Although I have presented a need for assistance with decision-making, when we examine the clinical domain, we can see that there is a need for more than just decision making. While at a simple level a drug-dose calculator may simply request some data about the patients weight, and apply a formula to produce a result, the problems I outlined above require more than this. In order for us to make, and justify, the decisions that we come to, we need more than just calculation. Instead, we need to be able to represent our domain, and our knowledge about that domain, and then use that, together with some patient data, to reason with this knowledge. Existing work on argumentation has given us different approaches to reasoning, in the form of different logical approaches, but I suggest that we need much more than a logic. Given this, we might ask about the use of argumentation as a Knowledge Representation (KR) formalism. There have been various definitions of what KR is, and we reproduce two here. Davis, Shrobe & Szolovits [23] defined KR as being:

- A surrogate for the real world
- A set of ontological commitments
- A theory of intelligent reasoning (both sanctioned and recommended inferences)
- A medium for efficient computation
- A medium for human expression

whilst Sowa [71] divides it into:

- Ontology
- Logic
- Computation

Given these, I think that existing work has concentrated on the latter three of the first group, and the latter two of the second. This suggests that there is a weakness in the current work on argumentation to sufficiently consider ontological issues. These criteria also they suggest if we are to develop a KR formalism based on argumentation, we need to address that while an inference mechanism (in our case, a logic) is clearly *necessary*, it is not *sufficient* for a KR system. This is an important point, as it allows us to start to make a clear break between what the logic must do (on which there is much debate in the argumentation field), what the KR formalism must do (on which there is much debate in the knowledge engineering field), and what the completed system must do (on which there is much debate in the clinical field). Despite all these debates, we may at least claim to have made some progress by assigning each problem to its correct area of debate. Furthermore, it allows us to describe exactly what the contribution of this thesis is: it provides a technique to link ontological committments with a theory of argument-based intelligent reasoning.

1.3.6 The Use of Strict Rules

Many of the defeasible-logic based formalisms include both strict and defeasible rules. For example, Prakken & Sartor give this example in their 1997 paper [62]:

Example 1.3.1. Let L be a propositional language where $a...d$ denote propositional formulae and $x...z$ are variables in a suitable meta-language. A rule, denoted r_n: $x \Rightarrow y$ denotes a defeasible rule, s, $\rightarrow$ a strict one and $\neg$ is strong negation

r_1: $a \Rightarrow x$ is married

r_2: $b \Rightarrow x$ is a bachelor

s_1: x is married $\rightarrow \neg$ x is a bachelor

s_2: x is a bachelor $\rightarrow \neg$ x is married

Similar examples are given by Garcia & Simari, for instance this one:

Example 1.3.2. Again using r, $\Rightarrow$ to denote a defeasible rule, s, $\rightarrow$ a strict one and $\neg$ is strong negation

r_1: bird(X) $\Rightarrow$ flies(X)

r_1: chicken(X) $\Rightarrow$ $\neg$flies(X)

s_1: penguin(X) $\rightarrow$ bird(X)

s_2: chicken(X) $\rightarrow$ bird(X)

s_3: penguin(X) $\rightarrow\neg$flies(X)

I would suggest that the use of strict rules is actually an attempt to address the problems caused by lack of ontological considerations. However, there are problems with this approach. Firstly, strict rules are used for a variety of different uses in the examples above. On the one hand, it solves the problem of relating (syntactically) unrelated terms, such as married and bachelor, to each other. On the other, it is used for capturing information about properties of sets of individuals (ontological information), such as penguins being birds.

The problem with this approach is that the strict rules in this, and other, examples, form part of a 'world model' - they tell us how terms are related to one another. This is very different information from the sort captured by the defeasible rules, and so should be separated out. Just as there are two main uses for strict rules, so there are two main objections. Firstly, for any large-scale application, with a large propositional language, the number of rules required to express negation between syntactically unrelated propositions becomes prohibitive. Secondly, in many domains, the ontological information comes from a different source than the defeasible information. Therefore, by separating the way in which we represent the different types of knowledge, we can separate the authoring of ontological information from the authoring of defeasible rules. This would make it easier to reuse the strict (ontological) information, as well as allowing different groups of authors to work on both types of information.

1.4 Uniting the clinical & technical problems

I have highlighted some of the problems of clinical care, as well as some technical weaknesses of existing argumentation systems. The intention is that we should be able to use an argumentation formalism as the basis for a medical DSS. If this is to be the case, we need to be clear on what we expect of it. Buchanan and Smith [17] have suggested that a medical DSS should:

1. Provide a solution at the same level of performance as a human expert
2. Use symbolic and heuristic reasoning rather than numeric and algorithmic procedures
3. Store knowledge separately from inference procedures
4. Provide explanations of their reasoning

Given the discussion above about the problems of knowledge, I would add that it must be capable of handling inconsistency. Using these criteria, and recalling that for now we loosely regard argumentation as a logic-based approach for constructing arguments for and against claims based on a body of knowledge, we can see that argumentation could provide (2), (3) and (4), as well as the resolution of inconsistency, and so would appear to be a good basis for a medical DSS. Specifically, (2) is satisfied by the logical asis for argumentation, rather than the numerical approaches of some other systems, (3) is satisfied by the separation between the knowledge stored (for example, in the type of defeasible rules above) and the way we reason with them, which is to construct arguments, and (4) is satisifed by the notion of the claim of an argument being based on the support of the argument. Although there is a body of theoretical work on argumentation, implementations have been much rarer, and generally not medically based. Most have concentrated on assisting human argumentation via the use of diagrams [66], [75], [46], rather than constructing a system that can reason. Given the definition of DSS above, we must concentrate our attention on Tallis/ProForma [72] and OSCAR [60] as these are the only implemented single-agent examples. Here, there is a small, but reasonably robust, body of evidence that an argumentation-based DSS can help improve clinician decision-making in prescribing [76], genetic risk assessment [44] and the diagnosis of breast cancer [59].

My interest is in one particular area. Given that we know that DSS can help improve care [45, 41], and that argumentation can and has been used to 'power' such systems, what are the properties of an argumentation system that we would expect if we wanted it to be useful for clinicians? I am not talking about the properties of a DSS - these might include a good user-interface, rapid and easy access, etc. Instead, if one were to build such a system, what would the underlying argumentation formalism look like? What would it need to do? What 'parts' would it need to have? We can perhaps answer this by returning to the two types of knowledge that are important in our domain. If we have a system that seeks to represent both types of knowledge, we may want to use different approaches for each type. On the one hand, we need a system that can represent static, non-contentious, value-free knowledge (that is considered clinically uninteresting) but on the other we want to represent rapidly changing, conflicting, value-sensitive knowledge, and use both of these to inform decision making. Arguments about therapy choice in post-surgical breast cancer frequently revolve around the differing outcomes of different treatments. Information on these is provided by the results of clinical trials, but there is also the background information that helps make sense of the trial results. Since this is the domain I am interested in, it seems sensible to suggest that the argumentation system should use the results of clinical trials to form arguments for and against various treatment options. The claims of arguments might be epistemic (e.g. Do I believe that someone has some disease, and if so, why) or based on action (e.g. Should I do X, and if so, why? And what are the likely consequences?). In addition, we might want it to acknowledge that different people place different weights on different outcomes, and will choose to differentially prefer certain pieces of evidence when coming to a conclusion. I also want to capture some of the richness of medical knowledge - for example, at one time I might want to make quite broad statements - 'This person should have chemotherapy' and at other times very precise statements - 'This person should have drug X at a particular dose schedule for 5 years'. I have already suggested that argumentation fits areas (2) and (4) of Buchanan and Smith's criteria. The first criterion, functioning at the same level as a human expert, is perhaps outside my control. Although previous work on argumentation has approached (3), the medical work to date has concentrated on modelling the recommendations of guidelines in argumentative form for decision making. However, the information on which these guidelines are based has not been represented in any detail, and so there has been a conflation of information (from the results of clinical trials) and the use to which it might be put (making decisions about treatment). In addition, work to date has generally failed to start with a defined set of predicates, instead assuming that they simply exist in some logical language . Although this works very well for small examples, if we want to build substantial applications, this approach seems more difficult. Therefore, in order to resolve these weaknesses, my new formalism should:

1. Model the results of clinical trials, and the background knowledge that provides the terms used to describe the results of the trials.
2. Model arguments for both belief and decisions - we have two main requirements: To know what might happen, and to know what we should do, and we need to be able to argue about both of these.
3. Represent different value-judgements in forming arguments: The move from knowledge about an intervention towards an argument for (or against) an intervention is dependent upon the integration of value judgements, and so we need to be able to represent these values, and the generation of arguments, in our system.
4. Take a piece of medical knowledge and use it to form different arguments: When we talk about using value judgements to help make arguments, we want to do so without having to simply repeat ourselves. We therefore want to find a way to have a single piece of knowledge, and use it in different ways.
5. Represent knowledge at different levels of abstraction: In some cases, our knowledge will be quite precise, and at others will be far more general, and the system needs to use both.

1.5 Outline of my argument-based solution

I will use a hybrid architecture to address this problem. It will consist of three areas:

- An Ontology
- A set of defeasible Rules,
- Arguments produced from the rules and the ontology

I will combine these approaches to deliver a KR system that is capable of producing arguments for and against both beliefs and intentions, and which will enable us to explore the impact of values on these arguments. This is a significant new piece of work, and draws on existing work in both the Description Logic/ semantic web and argumentation fields. As such, it should be of interest to those in both areas, as well as those working in Medical Informatics.

To model the results of clinical trials, I will represent the results of trials as rules, and the need for different types of arguments suggests that we will need different types of rules. This is described in chapter 3. The use of patient values to develop arguments about interventions, and the reuse of knowledge is described in chapter 4 and 5. To model the background knowledge in the domain I use an ontology, and this also allows the representation of knowledge at different levels of abstraction. This is described in chapter 3 and 4.

This work's novelty lies in the formalism it gives for linking a Description Logic (DL) ontology with an existing argumentation formalism, as well as giving definitions for certain types of rules and arguments, and an extensive case-study using over sixty defeasible rules developed from the literature. Description Logic researchers have only recently begun to address the issue of conflicting and inconsistent data, despite Tim Berners-Lee's identification of the issue over 6 years ago. This work will therefore be of interest to the Description Logic/ Semantic Web community, as it represents an advance in the use of defeasible reasoning in the context of the Semantic Web.

This work stands at the intersection of Argumentation, Knowledge Representation and Semantic Web technologies. However, these links are slightly unusual: It clearly has strong links to argumentation, and yet does not aim to produce substantial new formal results in argumentation; it draws some of its inspiration from Knowledge Representation, and yet will not produce a new KR formalism; and it uses some of the Semantic Web technologies, but generally supposes a single-agent, non-networked environment. Instead of concentrating on advancing each section in its area of strength, I intend to use them to address each other's blind spots. As such, this work should be of interest to workers in all three communities.

1.6 Contributions & Overview of Chapters

- This chapter has given an overview of some of the clinical and technical factors behind the work, as well some of the limitations of existing formalisms
- Chapter 2 reviews my chosen ontological formalism, and introduces the clinical ontology that I use throughout the rest of the thesis

- Chapter 3 presents an extension to an existing argumentation formalism. I take Garcia & Simari's DeLP system and introduce some new definitions to allow the incorporation of ontological knowledge. The resulting ontology argumentation formalism (OAF) is used in the rest of the thesis. The novel aspects of this chapter are the adapted and novel definitions together with some theoretical results

- Chapter 4 presents the work on incorporating values into the formation of arguments for action via the use of a rule-rewriting function. The novel work consists of the definitions and theoretical results about dialectical tree-formation

- Chapter 5 describes how arguments for actions can be used to hypothetically reason about possible consequences of actions. The novel work consists of definitions and proofs that I use to define criteria under which we can expect an OAF to generate ergonic and hypothetical arguments

- Chapter 6 presents a brief methodology for describing the sources of the rules and using them to decide on the preference status of arguments, as well as two algorithms for argument generation, and upper-bounds on their time-complexity

- Chapter 7 is a substantial case study, using rules developed from the references used in a breast cancer guideline. The arguments constructed from the rules are compared to the statements in the guidelines. The novel work consists of both the scale of the case study and the technique used to compare the formal arguments to the guideline.

- Chapter 8 summarises the main results of the thesis, and explores possible avenues for further work

1.6.1 Work Published from thesis

So far an earlier version of Chapter 3 has been published in the proceedings of the Nineteenth International Conference on Tools with Artificial Intelligence (ICTAI '08), and I am intending to submit further work for publication based on the results in chapters 4, 5 and 7. Currently, I am intending to submit the methodology and results of chapter 7 as a medical informatics paper, and an extended version of chapter 3 as a computer science paper on integrating argumentation and description logics.

Chapter 2

An Introduction to Argumentation & Ontology

This chapter introduces some of the existing work on argumentation and ontologies. It starts with a presentation of different argumentation formalisms developed over the last 15 years, and concludes by presenting the defeasible extended logic programming framework, DeLP. It then gives an informal introduction to the idea of an ontology, discusses a class of formalisms called Description Logics (DLs), and describes a breast cancer ontology written using a DL. I review some recent work on the logical and computational aspects of DLs, and I introduce a simple DL, $\mathcal{AL}$. I use a running example to demonstrate some of the techniques used for reasoning in DLs, and conclude by introducing a particular DL that I use in the rest of the thesis. The chapter introduces the breast cancer ontology, and discusses some of the practical problems in building such an ontology, as well as giving an overview of some of the available tools for working with such ontologies. This ontology then provides the predicates used in the defeasible rules in the next chapter.

2.1 Argumentation

Argumentation aims to reflect how humans use conflicting information to construct and analyse arguments, and central interests involve identifying arguments and counterarguments and evaluating them. In [19], Caminada and Amgoud suggest that argumentation is a reasoning process with four main steps:

1. Argument construction
2. Conflict detection
3. Determining the acceptability of arguments
4. Deciding on justified conclusions

There have been many different formalisms in argumentation, and a recent book [13] has described the major schools of work. In separate work, the authors have also shown [12] how an argumentative structure can be used to capture conflicts in the scientific literature.

2.1.1 Argumentation Formalisms

Besnard & Hunter [13] divide formal argumentation approaches into three main types:

1. Graph-based approaches

2. Defeasible-logic based approaches

3. Coherence-based approaches

I give a brief summary of an exemplary piece of work in each area below, before discussing one in more detail. I then go on to discuss some specific approaches to argumentation-based formalisations of the practical syllogism.

2.2 Argumentation Formalisms

2.2.1 Dung

In [27], Dung presents an abstract argumentation framework, consisting of set of arguments, A, and a set of attack relationships between the members of A, Attacks. An argumentation framework is a pair, $\langle A$, Attacks$\rangle$. The status of an individual argument depends on whether it is attacked by another argument, and whether the argument that attacks it is itself in turn attacked. The paper then defines semantics for the argumentation framework, depending on the attack relationships between arguments and sets of arguments.

The advantage of Dung's system is that its intial presentation is very clear, and it has sparked a great deal of related work (see [3, 6, 13] for examples). However, the crucial weakness for my purpose is that although Dung defines a way of resolving argument interactions, it starts from an atomic definition of an argument. The paper gives no definition of what an argument is, or how we might construct one from a knowledge-base. As a result, Dung's work provides an elegant approach for resolving conflicting arguments, but does not allow us to define and construct arguments. Since I am interested in using arguments to represent clinical knowledge, the structure and claims of these arguments, as well as their interaction, is significant. Although there has been some work embedding a system for constructing arguments in a Dung-style system for resolving conflicts [55], such approaches are not yet widely in use. There are also some underlying technical problems in trying to integrate two such different formalisms. for example, most argument formalisms that give a definition of an argument also contain the idea of a sub-argument, and often restrict the ways in which they can interact. However, as Dung's approach is abstract, and has no notion of sub-argument, Dung's argument semantics cannot easily take sub-arguments into consideration.

2.2.2 Prakken & Sartor

In [62], Prakken and Sartor introduce a defeasible-logic based argumentation system. Rules are conjunctions of literals in the body, with a single literal in the head, and may be either defeasible or strict in nature. An argument is the a sequence of defeasible and strict rules, the head of each rule in the sequence in considered to be a claim of the argument (hence arguments may have multiple claims).

Arguments are said to attack each other if they have two claims, ϕ and $\neg\phi$, or if they have claims such that the claim may be extended by the use of strict rules only to derive $\neg\phi$. For example, if we have one argument whose claim is ϕ and another whose claim is ψ, then we have no conflict, unless there is a strict rule of the form $\phi \rightarrow \neg\psi$ or $\psi \rightarrow \neg\phi$. Resolution of conflicting arguments is done on the basis of a preference order between the rules that make up the support of the argument, and this preference order is explicitly allowed to itself be defeasible - i.e. different proponents may differ in the preference ordering they use. The paper also gives sematics for the system, and proves certain important properties, such as completeness.

However, there are some technical weaknesses in their approach, as discussed in [19], specifically the fact that strict rules can allow for inconsistent conclusions to be drawn, all of which are considered justified. In addition, there is no implementation of the system available for use, and the authors do not explicity allow the use of rules with variables in as a schematic representation for multiple different ground rules, which later work does.

2.2.3 Besnard & Hunter

In [12], Besnard & Hunter define an argumetation system based on classical logic. An argument is a pair, $\langle\Phi,\alpha\rangle$, such that $\Phi \not\vdash \perp$, $\Phi \vdash\alpha$ and there is no $\Phi' \subset \Phi$ such that $\Phi' \vdash\alpha$, where $\vdash$ is deduction in classical logic. They then go on to define undercuts to arguments, where an undercut to $\langle\Phi,\alpha\rangle$ is an argument $\langle\Psi,\neg\beta\rangle$ and $\beta \in \Phi$. They then give definitions to allow us to consider only the most conservative of undercuts, and use these conservative undercuts to construct a tree of arguments where each argument is an undercut to the node above it in the tree. The work then goes on to define ways in which the interaction of arguments can be assessed, in particular by considering the extent to which arguments contradict, and to pries for pruning and compressing argument trees in order to make them more tractable.

For my purposes, there are several weaknesses. Firstly, the presentation is very abstract, and as a result there is a gap between the description give in the paper and an obvious understanding of how it could be used for actual reasoning. Secondly, the use of classical logic necessitates more complex reasoning than that the use of defeasible logic. Therefore, although the work is interesting, it is not as easy to use as the basis for my work as some other proposals.

2.3 Argumentation Systems

I am interested in using argumentation to capture the knowledge in clinical trials, and that I want to extend an existing argumentation formalism. Graph-based approaches are not appropriate as a basis for the thesis, as they contain no definition of what constitutes an argument. Coherence-based approaches could provide a basis for the work, but have been based on first-order logic (FOL) and there are no implemented systems. This leads us to the defeasible-logic based approaches. Of these, the Defeasible Extended Logic Programming (DeLP) [32] formalism by Garcia & Simari is attractive as it is described in a single, precise, self-contained publication, with an implementation available online. In addition, it resolves some of the technical problems of earlier approaches such as [62].

2.4 DeLP

Garcia & Simari

I do not give a complete review of DeLP here, as chapter 3 is based on it, and there would be a great deal of repetition. In summary, DeLP has two main aspects: argument construction and the resolution of conflicting arguments. As with other defeasible logic systems, DeLP has rules with conjunctions of literals in the body and single literals in the head, and rules may be either strict ($a \rightarrow b$) or defeasible ($a \Rightarrow b$). One minor development is the use of schematic rules to stand for multiple ground rules: Thus instead of $Bird(Tweety) \Rightarrow Flies(Tweety)$, DeLP allows us to write $Bird(X) \Rightarrow Fly(X)$ as a defeasible rule and supply $Bird(Tweety)$ as a fact. Argument construction is similar to that in other defeasible logic systems, in that an argument is of the form $\langle A, \phi \rangle$, where A is a set of rules, such that when considered in conjunction with a set of facts, A is a minimal set of consistent rules that provides a derivation for ϕ. Interaction between conflicting arguments is resolved by constructing a tree of arguments (a so-called dialectical tree), where the defeaters of an argument are used as its leaf nodes. The root argument is then considered warranted if all of its defeaters are defeated. Definitions and examples of this are given in chapter 3. This approach is attractive as it provides a concise way of encompassing the conflicts and interactions between arguments. Fundamentally, the defeat status of an argument is based on a preference-based approach to resolving conflict, based on the relative weight of the rules in the support of the arguments.

Compared to other argumentative approaches, DeLP has several advantages. Firstly, unlike the graph-based approaches, DeLP, like other defeasible-logic approaches provides a practical basis for representing and reasoning with knowledge. Unlike the coherence-based approaches, the inferential aspects of the system are very simple, consisting of a defeasible form of modus ponens reasoning alone, and hence reasonably imple to implement. When compared to other defeasible-logic approaches DeLP solves more of the problems suggested in [19], and the original presentation has some simple but useful aspects, such as the use of schematic rules, which makes presenting rules far easier.

2.5 The Practical Syllogism

Considering reasons for acting has been a central concern of philosophy, and later computer science, for thousands of years. Aristolte is often credited with first presenting practical reasoning as a syllogism, and his overall scheme remains in use. A commonly accepted approach would be to say that if in some circumstance A, action B will achieve end X, and I believe X to be desirable, useful or good, then I should perform action B. Although there have been numerous formalisms of the practical syllogism, there have been some recent argumentation-based formalisms that are of particular relevance given the work in chapters 4 and 5, and I will briefly review these here.

One of the novel aspects of OAF is the use of an ontology. Since the other formalisms have no equivalent to the ontology in OAF, I will regard the ontology as analogous to a set of strict rules, as used in DeLP or Prakken & Sartor's work. Compared to the work on epistemic argumentation, there has been less work on argumentation for decision-making. There are large amounts of other work on different aspects of the practical syllogism [78, 34, 8], and Atkinson's work on the practical syllogism and critical questions [8] has also been used as the basis for other work on practical reasoning [20, 73], but my interest here is in comparing the work in the thesis with similar predominantly logical approaches. The most notable work includes the Logic of Argumentation (LA) [50], and work by Amgoud [3, 6, 4], and Bench-Capon [10]. For our purposes LA is relatively uninteresting; Although used to generate arguments for decisions, including actions, the relationship between the situation and the action is explicity stated in the rules that LA uses. Hence LA can be used to generate arguments for (resp. against) different courses of action, but it does not provide a formalism of the practical syllogism, and it is the other two formalisms I will examine here.

2.5.0.1 Amgoud

Amgoud [4, 3] has presented a 'unified' framework for inference and decision as an extension of previous work on argumentation frameworks [5] which uses a mixture of strict and defeasible rules in some logical language $\mathcal{L}$ which is closed under Negation, and four sets of formulae in $\mathcal{L}$, Decisions, Beliefs, and Positive and Negative goals. Arguments are divided into three types, where A_1.... A_n are arguments or formulae in $\mathcal{L}$, ψ are formulae in $\mathcal{L}$ and d is a member of the set of decisions.

1. Epistemic, of the form: A_1.... $A_n \to \psi$
2. Recommending, of the form: A_1.... $A_n \to d$
3. Decision, of the form: A_1.... A_n, $d \to \psi$

Epistemic arguments are similar to those in other formalisms (including mine), and will not be further considered; Recommending arguments are those that use beliefs to make arguments for/ against decisions; Decision arguments are those that link epistemic arguments and beliefs with decisions to form new beliefs. The work contains several results of interest, namely that all proper subarguments of recommending or decision arguments are epistemic, and defines an argument comparison based on the strength of arguments supporting each decision.

2.5.0.2 Bench - Capon

Bench- Capon and Prakken (BCP) present a framework that uses a propositional modal logic as a language for strict and defeasible rules to form arguments via modus ponens, and uses the modal operator D in conjunction with propositional literals to denote goals (thus Dψ to mark that $\psi$is a desired goal and $\neg$Dϕ to mark that ϕ is an undesired goal, with the identity $\neg$D $\neg\psi \equiv$ Dψ. Furthermore, they differentiate between formulae that are controllable and uncontrollable where only controllable formulae may participate in the following relationship:

$$\frac{\psi \Rightarrow \phi \;\; D\phi}{D\psi}$$

Their use of a single modal operator allows for chaining of desirable goals, as the following example shows:

Example 2.5.1. Given the following rules:

$$a \wedge b \Rightarrow c$$
$$d \wedge e \Rightarrow b$$

and the goal-base

$G = \{Dc\}$

with b, e being choosable, then we can form an argument for Dc

$$\frac{a \wedge b \Rightarrow c \;\; Dc}{Db}$$

and on the basis of this:

$$\frac{d \wedge e \Rightarrow b \;\; Db}{De}$$

This example shows an important aspect of BCP's framework, namely the way in which the desirability of of some formula ϕ can lead to an argument for the desirability of some other formula ψ via a method analogous to that of backward chaining. This is an interesting approach, and certainly has some intuition behind it - if b leads to c, and e leads to b, and we desire c, then we should desire b, and therefore desire e.

The above summary is designed to allow the body of the thesis to be set in context, particularly the work in chapter 3, 4 and 5.

2.6 Ontological Introduction

This chapter provides the ontological background for the thesis. I must start by saying that I am a consumer of ontologies; that is, this thesis is not concerned with developing new ideas in ontologies, or solving problems with existing ones. In addition, the presentation here is necessarily abbreviated, and is heavily skewed towards covering those areas that are of interest in my work; for a much more comprehensive analysis, the reader is directed to [9].

2.6.1 Informal Presentation

We should start by distinguishing the philosophical notion of ontology (the enquiry into what exists and why) and applied ontology. Perhaps the best known definition of ontology (in the second sense) is:

> "An ontology is a specification of a conceptualization" [35]

The 'conceptualization' here comes from [33], where a conceptualization is 'an abstract simplified view of the world'. Even more simply, we may view an ontology as a *model* of what exists in the world. Traditionally, ontologies have contained *Instances*, which are grouped into *Classes* (a class is a set of similar instances) and *Properties* (a relationship between some instances). It is normally assumed that any member of a sub-class is also a member of the super-class, as the following example shows:

Example 2.6.1. Imagine an ontology with the classes People, Women and Men. We know that all Men and Women are People, and we might say that no individual can be both a man and a woman. We would say that Women and Men are sub-classes of People, and we might define the classes Men and Women as being disjoint with each other.

As the example above shows, because ontological terms often refer to things in the world (e.g. People), I distinguish the ontological terms by use of a sans-serif font (People). Because classes (potentially) refer to many individuals, it is often customary to use plural nouns (hence People not Person). In addition, we can arrange the classes (and sometimes the properties) in a hierarchy, so that the sub-classes (lower down the tree) derive from the super-classes.

2.6.2 Historical Background

The practice of constructing ontologies is not new, and is often traced to Aristotle's use of the word *Category* and the medieval developments of *Porphyry's Tree* (itself a margin note in a commentary on Aristotle) [71]. Since the 1960s there has been work on semantic nets [21], Brachman and Levesque's work on logical bases for frames and Sowa's work on Conceptual Graphs [71]. Although the field may have started with a strong logical foundation, some of the intervening work was more interested in capturing facts about the world than in ensuring its underlying logical correctness. More recent work, such as that by Brachman and Levesque [64], has concentrated on trying to maintain the ease of use of semantic networks whilst also supplying formal semantics. Such work has culminated in the the recent development of Description Logics and languages such as KIF [52] and OWL [58]. However, all of the approaches contain some form of inheritance between classes, in that if one class of instances is a sub-class of another, every member of the sub-class is also a member of the superclass.

2.6.3 Uses & Examples

Over the last few years, there have been several attempts to build and use ontologies for practical applications in the biomedical field. Two of these are the Gene Ontology and SNOMED-CT. The Gene Ontology was driven by the desire to annotate genomic data in a standardised fashion. Although originally developed using a bespoke format, it is now available as an OWL file. In addition to the core genomic ontology which contains over 21 000 terms, it is also an associated ontology to describe the evidence supporting each statement. SNOMED-CT has approximately 350 000 terms, and was developed by the American College of Pathologists originally as a controlled vocabulary for describing illnesses and treatments, but has evolved into a DL-based ontology. It is the proposed ontology for use in the NHS current 'Connecting for Health' IT project, where, in conjunction with the HL7 messaging standard it is supposed to provide the basis for the electronic patient record. It is also available as an OWL file.

2.6.4 Recent Work

Recent developments in applied ontology have increased its importance and visibility. The first is the development of formal semantics, provided by DL, along with a description of the different computational complexity implications of different formalisms. The second was the development of formats, such as OWL. I shall discuss both logics and format below, and I shall assume that the reader is familiar with standard first-order logic.

2.7 Description Logics

2.7.1 Syntax

DLs are typically written in a 'variable-free' syntax. For example, in first-order logic (FOL), the predicate $Men(x)$ might be used to denote the set of all men and we might write $Men(Matt)$ to denote that $Matt$ is a member of this set (and hence a man). A DL term to denote the same set would be Men and to denote membership, Men(Matt). This syntax extends to binary predicates, as well. So to express the fact that a married man is someone who has a wife, we might write (in FOL):

$$\forall x, \exists y.\ MarriedMen(x) \equiv Men(x) \wedge hasSpouse(x, y) \wedge Women(y)$$

whereas, following the DL syntactical conventions, we would write:

$$\mathsf{MarriedMen} \equiv \mathsf{Men} \sqcap \exists \mathsf{hasSpouse.Women}$$

2.7.2 Basic DL

DL are a fragment of first-order logic, restricted to monadic predicates (Concept Names), binary predicates (Role Names) and constants (Individuals). New concept names are defined recursively using certain (restricted) combinations of concept and role names. Due to these restrictions, most description logics have good computational properties in terms of decidability and computational complexity. However, there is no agreement on a single 'perfect' DL, and so different description logics have developed. These vary in the restrictions they impose on the way that concepts and roles can be combined to form new definitions, and the implications of these restrictions for the complexity of the logic (e.g. the move from the $\mathcal{SHIF}(\mathcal{D})$ DL to the $\mathcal{SHOIQ}$ DL results in an increase in complexity [38]; for the naming scheme, see 2.7.3.3). When we come to implement systems that use such logics, this information allows us to tune the choice of formalism to the application and predict its performance.

Because description logics are a family of different logics, I have taken a two-step approach to their introduction. The work in this thesis is relatively indifferent to which DL is used, and so I start by discussing some basic terminology such as entailment and subsumption checking with reference to a very simple DL, the 'attributive language' ($\mathcal{AL}$) logic [9]. I then go on to discuss some more complex topics solely in the context of my chosen DL.

One of the features of DL is the emphasis that they place on recursively defining (non-ground) formulae so as to develop predicates with complex meanings. Much of the work on authoring and reasoning with DL has concentrated on such formulae, and a useful distinction can be made between the T-Box and the A-box of an ontology, which I denote $\mathcal{K}^{\mathcal{T}}$ and $\mathcal{K}^{\mathcal{A}}$ for any given DL vocabulary $\mathcal{K}$. The T-Box contains non-ground formulae (the terminology), while the A-Box contains ground formulae (assertions) about individuals. As we shall see, T-box reasoning can itself become quite complex, and indeed A-Box reasoning has in some ways been neglected in the literature. This is in contrast to (for example) databases, where the focus is on instance data, rather than schema definition.

2.7.3 DL Formalisms

Although all DL share a common set of assumptions, because there are many different DL, the precise definitions may vary according to the logic. However, I shall consider the $\mathcal{AL}$ logic as an exemplary DL in the section below.

Definition 2.7.1. Let $\mathcal{K}$ be a simple DL vocabulary. Then $\mathcal{K}$ consists of four sets :

- C, the set of Concept names
- R, the set of Role names
- I, the set of Individuals

- $\#$, the set of logical symbols

where Concept names are monadic predicates, Role names are binary predicates, Individuals are constants in the language and $\#$ = {¬, ⊓, ⊔}

I start by describing the set of ground literals, or assertions: the A-Box.

2.7.3.1 The A-Box

The A Box, A, is a finite set of assertions. We represent these assertions as:

Women(Mary)
Men(Matt)
hasChild(Mary, Matt)

which expresses the fact that Mary is a Woman, Matt is a Man, and Mary is related to Matt via the binary hasChild predicate. Mary is referred to as an *instance* of Women, and Matt is an *instance* of Men, and since they are named, we refer to them as ***concrete*** individuals. Since Women is a concept name, and hasChild is a role name, we refer to Women(Mary) as a concept assertion and hasChild(Mary, Matt) as a role assertion, and hence atoms and literals are the assertions (or the 'A-Box') of a DL. The elements of the vocabulary are used to define atoms and literals in the $\mathcal{AL}$ DL.

Definition 2.7.2. Let $\mathcal{K}$ be a vocabulary as above. For some concept name $c \in C$, role name $r \in R$ and individual $i_1, i_2 \in I$, an $\mathcal{AL}$ ***atom*** is either $c(i_1)$ or $r(i_1, i_2)$ A ***literal*** is an atom or a negated monadic atom, and the set of all literals in $\mathcal{AL}$ for a vocabulary $\mathcal{K}$ is denoted $\mathcal{K}^{\mathcal{A}}$.

Example 2.7.3. Let $\mathcal{K}$ be some vocabulary. Then if Men $\in C$,hasCat $\in R$, and Matt, Puss_in_Boots $\in I$, Men(Matt) and hasCat(Matt, Puss_in_Boots) are atoms in $\mathcal{AL}$, and ¬Men(Matt) is a literal but ¬hasCat(Matt, Puss_in_Boots) is not.

I now move on to describe non-ground formulae, the T-Box.

2.7.3.2 The T-Box

The simple elements of the T-Box are atomic concepts (denoted here by A,B) and atomic roles (denoted here by P,Q). The $\mathcal{AL}$ DL has T-Box formulae defined according to the following rules:

Definition 2.7.4. Let $\mathcal{K}$ be a vocabulary as above. For some concept name $c, d \in C$, role name $p, q \in R$, the following are formulae in the $\mathcal{AL}$ description logic:

$\top$ (universal concept)

$\bot$ (bottom concept)

c (atomic concept)

$\neg c$ (atomic negation)

$c \sqcap d \sqcap ... \sqcap$ C(intersection)

$c \sqsubseteq d$ (subclass)

$\forall p.c$ (value restriction)

$\exists p.\top$ (limited existential quantification)

$c \equiv d$ (equivalence)

These restrictions are referred to as the allowable ***constructors*** for $\mathcal{AL}$, and the set of all such formulae for some vocabulary $\mathcal{K}$ if referred to as $\mathcal{K}^{\mathcal{T}}$.

The first four choices are hopefully familiar; $\top$ is the 'top' concept, where for any constant c in the ontology, $\top(c)$ is true, and $\perp$ is bottom, such that for any constant c, $\perp(c)$ is false. For some atomic concept A, $\neg A$ is its negation. Note that negation in DLs is strong (classical) negation, rather than negation as failure. For any two atomic concepts A and B, $A \sqcap B$ is the intersection of them. $\forall$R.A and $\exists$R. $\top$ impose restriction on binary properties (roles). $\forall$R.A restricts the range of a role to be of the class A, while $\exists$R. $\top$ allows one to assert the existence of some constant d such that the relationship $R(c, d)$ holds; however, there are no restrictions on which concept d is a member of. These last two are best explained through example.

Example 2.7.5. The following are examples of some well formed formulae in $\mathcal{AL}$:

$\forall p.a$ allows us to restrict the range associated with a role to a concept; thus if we had a role, hasChild, we might say that all the individual fillers of that role must be Human, which we would write as $\forall$hasChild.Human. As a result, if there are two individuals who stand related by the role hasChild (such as hasChild(Frank,Bob)), then we are able to infer that Human(Bob) is true.

$\exists p.\top$ allows us to define the fact that there must exist at least one individual for whom the relationship holds; however, we do not define the concept to which this individual belongs. Thus Married $\equiv\exists$hasSpouse. $\top$ allows us to express the fact that all married people must have at least one spouse; it does not, however, allow us to restrict that concept to which the spouse belongs. Therefore hasSpouse(Bob,Puss_in_Boots) would be a valid formula in the A-box, and would allow us to infer Married(Bob).

Note that the $\mathcal{AL}$ logic is a relatively simple one. For example, it does not allow unions of classes, or negation of non-atomic concepts. Thus for two atomic concepts c, d in $\mathcal{AL}$, neither $c \sqcup d$ nor $\neg(c \sqcap d)$ are valid formulae. Again, more expressive DLs allow such formulae.

Now we have definitions for $\mathcal{K}^{\mathcal{T}}$ and $\mathcal{K}^{\mathcal{A}}$, we can define other elements of the language.

Definition 2.7.6. For some vocabulary $\mathcal{K}$, $\mathcal{O}_{\mathcal{K}}$ is the set of ontological formulae in $\mathcal{AL}$, where $\mathcal{O}_{\mathcal{K}} = \mathcal{K}^{\mathcal{T}} \cup \mathcal{K}^{\mathcal{A}}$.

Therefore, $\mathcal{O}_{\mathcal{K}}$ is the set of all formulae from $\mathcal{K}$ that fulfill the criteria for being either a literal in the A-Box or a formula in the T-Box in $\mathcal{AL}$. Since DL are a subset of FOL, we would expect them to contain ideas of consistency and satisfiability, and definitions for these are given below. However, a particular set of consistent, satisfiable formulae is given a particular name, which it is useful to define here.

Definition 2.7.7. Let $\mathcal{K}$ be some vocabulary and $\mathcal{O}_{\mathcal{K}}$ the set of $\mathcal{AL}$ formulae. Then an $\mathcal{AL}$ ***ontology***, $\Delta \subseteq \mathcal{O}_{\mathcal{K}}$ is a consistent, satisfiable set of ontological formulae. We denote the A-Box of the ontology, $\mathcal{K}^{\mathcal{A}} \cap \Delta$ as Δ^{A}, and the T-Box of the ontology, $\mathcal{K}^{\mathcal{T}} \cap \Delta$ as Δ^{T}.

2.7.3.3 Defining new Concepts

It may be the case that the T-Box of an ontology, Δ^T consists simply of some set of concept names. Often, however, we will add new concept names by defining them in terms of existing concept and role names in accordance with the restrictions described above. To do this, we define new axioms in the T-Box of an ontology in terms of existing definitions. New definitions are introduced through the use of axioms.

Definition 2.7.8. Let $\mathcal{K}$ be some vocabulary, C the set of concept names in $\mathcal{K}$ and $\mathcal{K}^{\mathcal{T}}$ be the set of non-ground formulae from $\mathcal{K}$ for the $\mathcal{AL}$ logic. Then an axiom is of one of the following forms:

A $\sqsubseteq$ B

A $\equiv \phi$

where A, B $\in C$ are concept names and $\phi \in \mathcal{K}^{\mathcal{T}}$ is a formula in $\mathcal{AL}$ as defined previously. A $\equiv$ B is read as 'A is equivalent to B' and A $\sqsubseteq$ B is read as 'A is a subconcept of B'. Using such axioms, we can then recursively define new concepts.

Example 2.7.9. Consider some T-Box Δ^T such that Δ^T ={People, Women}. Then we could add a new concept, Men, and use an axiom to define it as so:

Men $\equiv$ People $\sqcap \neg$Women

Note that People $\sqcap \neg$Women is a valid formula given the constructors in $\mathcal{AL}$.

Although we can build up definitions via such axioms, it is still useful to distinguish those concepts that are defined by an axiom and those that are not.

Definition 2.7.10. In any T-box Δ^T, we divide the concepts into name symbols, N^Tthat occur on the lefthand side of any axiom and base symbols, B^T that occur only on the righthand side.

Example 2.7.11. Consider some T-Box Δ^T such that:

Men $\equiv$ People $\sqcap \neg$Women

Mothers $\equiv$ Women $\sqcap \exists$hasChild.$\top$

Fathers $\equiv$ Men $\sqcap \exists$hasChild.$\top$

Parents $\equiv$ People $\sqcap \exists$hasChild.$\top$

then N^T ={Men, Mothers, Fathers, Parents} and B^T ={People, Women, hasChild}.

The example above illustrates the way in which recursive definitions allow increasingly complex formulae in the ontology, as well as showing a limitation of $\mathcal{AL}$. Given the definitions above, it may seem more sensible to write:

$\mathsf{Parents} \equiv \mathsf{Mothers} \sqcup \mathsf{Fathers}$

But we recall that $\sqcup$ is not one of the constructors in $\mathcal{AL}$, as mentioned above, and therefore $\mathsf{Parents} \equiv \mathsf{Mothers} \sqcup \mathsf{Fathers}$ would not be a formulae in $\mathcal{K}^{\mathcal{T}}$ for $\mathcal{AL}$.

Named formulae (that is, those elements of $\mathcal{K}^{\mathcal{T}}$ that are in some ontology Δ and are in $N^T \subset \Delta$) are significant, as they introduce individuals whose existence is inferred, through an assertion of membership of some concept in N^T, rather than stated in the A-Box.

Example 2.7.12. Consider the T-Box above. Then if we have an assertion in the A-Box, Fathers(Bob) we can consider the T-Box axiom $\mathsf{Parents} \equiv \mathsf{People} \sqcap \exists \mathsf{hasChild}.\top$ and from that we may infer that there must be some individual, i, such that $\top(i)$ and hasChild(Bob,i) both hold.

Individuals such as i are known as ***abstract*** individuals, as they are not explicitly introduced and named in the A-Box. They may occur because of the existential quantification allowed in DL formulae, and such DL are decideable under certain restrictions. The distinction between concrete and abstract individuals is of importance for work in both this and the next chapter.

A note on naming:

Because certain patterns of restrictions in acceptable constructors recur, there is a convention that relates the naming of the logic to these restrictions. The DLs that are defined by these restrictions are referred to by acronyms based on the constructors allowed in T.

Abbreviation	Constructor
$\mathcal{AL}$	Constructors as defined in 2.7.3
$\mathcal{C}$	Complementation ($\neg$) of arbitrary concepts
$_{R^+}$	Transitive Roles
$\mathcal{S}$	Shorthand for$\mathcal{ALC}_{R^+}$
$\mathcal{U}$	Union
$\mathcal{I}$	Inverse Roles
$\mathcal{O}$	Nominals (one of a set)
$\mathcal{H}$	Role Hierarchies
$\mathcal{D}$	Datatype Roles (Roles with datatype literals as the role filler)
$\mathcal{F}$	Functional roles (of the form $\leq$1 R.C)
$\mathcal{N}$	(unqualified) number restrictions
$\mathcal{Q}$	Qualified number restrictions

Table 2.1: Naming convention for more complex Description Logics based on the DL constructors they incorporate. Note that $\mathcal{S}$ is used as a shorthand for $\mathcal{ALC}_{R^+}$.

Thus, a DL with the name $\mathcal{SHIF}$ is $\mathcal{ALC}_{R^+}$ extended with role hierarchies ($\mathcal{H}$), inverse roles ($\mathcal{I}$)and functional role restrictions ($\mathcal{F}$). $\mathcal{SHIF}$ is important as it is the basis for OWL-Lite, while $\mathcal{SHOIN}(\mathcal{D})$ is the basis for OWL-DL. As a result, OWL-DL allows considerably more complex formulae than OWL-Lite.

This concludes the brief introduction to DL syntax. I now go on to consider the semantics for DLs, before returning to the more complex logics in the next section.

2.7.4 Semantics

The semantics of a DL are given in a similar way to those of FOL, but with the simplification that we only need consider unary and binary predicates. We start with an interpretation for the atomic concepts and roles in the set of base terms (the elements of B^T and literals that contain terms in B^T), and then inductively extend this definition to cover the other concepts, roles and individuals.

Definition 2.7.13. For some ontology, Δ, an interpretation, $\mathcal{I}$ is a non-empty set $\Delta^{\mathcal{I}}$, referred to as the domain of the interpretation. The interpretation function assigns to every atomic concept A a set $\mathcal{A}^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$ and to every atomic role R a binary relation $\mathcal{R}^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$.

The interpretation function is then extended to concepts in the set of name terms by the following inductive definitions:

Definition 2.7.14. For some vocabulary $\mathcal{K}$, the interpretation of formulae in $\mathcal{K}^{\mathcal{T}}$ for the $\mathcal{AL}$ logic is below:

$\top^{\mathcal{I}} = \Delta^{\mathcal{I}}$

$\bot^{\mathcal{I}} = \emptyset$

$(\neg A)^{\mathcal{I}} = \Delta^{\mathcal{I}} \setminus A^{\mathcal{I}}$

$(C \sqcap D)^{\mathcal{I}} = C^{\mathcal{I}} \cap D^{\mathcal{I}}$

$(\forall R.C)^{\mathcal{I}} = \{a \in \Delta^{\mathcal{I}} | \, (\forall b.(a,b) \in R^{\mathcal{I}} \to b \in C^{\mathcal{I}})\}$

$(\forall R.\top^{\mathcal{I}}) = \{a \in \Delta^{\mathcal{I}} | \, \exists b.(a,b) \in R^{\mathcal{I}}\}$

Once we have defined interpretations for all the elements of the T-box, we define what it means for concepts to be equivalent:

Definition 2.7.15. For two concepts, C, D $\in T$, C and D are equivalent if all their interpretations are equivalent. We write C $\equiv$ D iff $C^{\mathcal{I}} = D^{\mathcal{I}}$ for all interpretations $\mathcal{I}$.

So far we have been dealing with assigning interpretations to the T-box, and we now extend this to individual names (i.e. to ground formulae in $\mathcal{K}^{\mathcal{A}}$).

Definition 2.7.16. An interpretation $\mathcal{I} = (\Delta^{\mathcal{I}})$ maps atomic concepts and roles to sets and relations, and also maps individual names a to an element $a^{\mathcal{I}} \in \Delta^{\mathcal{I}}$. The unique names assumption is interpreted as if a, b are distinct names, then $a^{\mathcal{I}} \neq b^{\mathcal{I}}$. The interpretation $\mathcal{I}$ satisfies the concept assertion C(a) iff $a^{\mathcal{I}} \in C^{\mathcal{I}}$ and it satisfies the role assertion R(a,b) if $(a^{\mathcal{I}}, b^{\mathcal{I}}) \in R^{\mathcal{I}}$.

An interpretation satisfies the ABox, A, if it satisfies each assertion in A. In this case, we say that $\mathcal{I}$ is a model of A. $\mathcal{I}$ satisfies an assertion α or an ABox A with respect to a T-Box T if in addition to being a model of α or A, it is a model of T.

Note that I have not dealt with the interpretation of the constructors found in other DLs. For details of how these constructors are interpreted, see [36].

2.7.5 Inference

Inference in Description Logics allows us to make new assertions on the basis of existing ones; for example, in the small ABox above, from Women(Mary) we should be able to infer that Mary is also member of People. We define some formal properties below:

1. **Satisfiability** A concept C is satisfiable with respect to T if there exists a model $\mathcal{I}$ of T such that $C^{\mathcal{I}}$ is non-empty. In this case we also say that $\mathcal{I}$ is a model of C.

2. **Subsumption** A concept C is subsumed by a concept D with respect to T if $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$ for every model $\mathcal{I}$ of T. In this case we write C $\sqsubseteq_T$ D or $T \models$ C $\sqsubseteq$ D.

3. **Equivalence** A concept C is equivalent to a concept D with respect to T if $C^{\mathcal{I}} = D^{\mathcal{I}}$ for every model $\mathcal{I}$ of T. In this case we write C $\equiv_T$ D or $T \models$ C $\equiv$ D.

4. **Disjointness** Two concepts C and D are disjoint with respect to T if $C^{\mathcal{I}} \cap D^{\mathcal{I}} = \emptyset$ for every model $\mathcal{I}$ of T.

If the T-Box T is clear from the context, we sometimes drop the qualification "with respect to T". It turns out that the crucial aspect is checking the satisfiability of concepts, for many other inferences can be reduced to (un)satisfiability [37]. In order to see how we can use DLs, let us look at a small running example to illustrate the definitions for satisfiability, consistency checking and entailment.

Example 2.7.17. Consider an $\mathcal{AL}$ ontology Δ, such that Δ^T is:

$$\mathsf{Men} \equiv \mathsf{People} \sqcap \neg\mathsf{Women}$$
$$\mathsf{Mothers} \equiv \mathsf{Women} \sqcap \mathsf{hasChild}.\top$$
$$\mathsf{Fathers} \equiv \mathsf{Men} \sqcap \mathsf{hasChild}.\top$$
$$\mathsf{Parents} \equiv \mathsf{People} \sqcap \mathsf{hasChild}.\top$$

and Δ^A is:

Men(Matt)

I shall use this small ontology in the examples below.

To start, we expand the terms of the T-Box so that we can see how they relate to one another.

2.7.5.1 Eliminating the T-Box

Recall the division of the T-Box into name symbols, $\mathcal{N}_T$ and base symbols $\mathcal{B}_T$, where for some concept definition A $\equiv$ D such that D contains only base symbols. Now for each concept C we define the expansion of C with respect to T as the concept C' that is obtained by replacing each occurrence of a name symbol A in C by the concept D, where A $\equiv$ D is the definition of A in T', the expansion of T. It follows from this expansion that C $\equiv_T$ C', and from this we deduce that C is satisfiable w.r.t. T iff C' is satisfiable. Similarly, we can show show that C and D are disjoint w.r.t. T iff C' and D' disjoint.

Example 2.7.18. Continuing from the example above, to check if Women and Men are disjoint, we calculate the expansion of the concept:

$$\mathsf{Women} \sqcap \mathsf{Men}$$

with respect the T-Box described which results in the concept:

$$\mathsf{Women} \sqcap \mathsf{People} \sqcap \neg\mathsf{Women}$$

Since this formula contains the intersection of Women and ¬Women, it is clearly unsatisfiable. Therefore we may conclude that Women and Men are disjoint.

In general, when we talk of formulae being disjoint, we need to do so with respect to the T-Box of a specific ontology, Δ. However, the process of T-Box elimination removes this caveat, as the example above shows.

2.7.5.2 ABox Reasoning

We can give a formal definition of consistency and entailment to the A-Box in a method similar to that for the T-Box. We simply say that A is consistent if it is consistent with respect to the empty T-Box, and is consistent with respect to a T-Box T if there is an interpretation that is a model of both A and T.

Example 2.7.19. Consider some ontology Δ, such that:

Δ^A= {Men(Matt), Women(Matt)}

Then Δ^A is consistent (with respect to the empty T-Box), because without any further restrictions on the interpretation of Men and Women, the two elements have been interpreted in such a way that they have a common element. In the same way as we expand the T-Box for concepts, consistency checking can be reduced to checking an expanded ABox. We define the expansion of A with respect to T as the ABox A' by replacing each concept and role assertion C(a), R(a,b) by the assertion C'(a), R'(a,b), where C', R' are the expansions of C, R with respect to T. We say that A is consistent w.r.t. T iff its expansion A' is consistent.

Example 2.7.20. Continuing the example above, if Δ^T is defined as:

$$\mathsf{Men} \equiv \mathsf{People} \sqcap \neg\mathsf{Women}$$

$$\mathsf{Mothers} \equiv \mathsf{Women} \sqcap \mathsf{hasChild}.\top$$

$$\mathsf{Fathers} \equiv \mathsf{Men} \sqcap \mathsf{hasChild}.\top$$

$$\mathsf{Parents} \equiv \mathsf{People} \sqcap \mathsf{hasChild}.\top$$

and

Δ^A= {Men(Matt), Women(Matt)}

Then Δ^A is not consistent with respect to the T-Box.

Continuing the example above (2.7.20), we have the ground formula, Women(Matt) ⊓ Men(Matt). If we expand this A-box, we have the expanded formula Women(Matt) ⊓ People(Matt) ⊓ ¬Women(Matt). Clearly this is inconsistent.

We can apply a similar expansion of terms to T-Box reasoning, and T-Box entailment can then be reduced to unsatisfiability [38], as shown below:

Definition 2.7.21. For concepts $C, D \in \Delta^T$:

1. C is subsumed by D iff $C \sqcap \neg D$ is unsatisfiable
2. C and D are equivalent iff both ($C \sqcap \neg D$) and ($D \sqcap \neg C$) are unsatisfiable
3. C and D are disjoint if $C \sqcap D$ is unsatisfiable

In the next chapter, we shall see how we use ontological entailment as part of defeasible reasoning. Since ontological entailment is well-defined for DLs, we denote it through a simple relationship.

Definition 2.7.22. For a vocabulary $\mathcal{K}$ and ontology Δ, we say an assertion $\alpha \in \mathcal{K}^{\mathcal{A}}$ is ***entailed*** by an A-Box Δ^A with respect to a Δ^T if every model of Δ^A wrt Δ^T satisfies α. We denote this $\Delta \vdash_{Ont} \alpha$.

2.7.5.3 Ontologies, FOL and Databases

DLs are often compared to databases, with Δ^T playing the part of the database schema and Δ^A the individual row entries. However, DLs are based on FOL, and depend on Open World semantics. For example, if the only assertion about Peter is hasChild(Peter,Tony), then in a database this would normally be understood as the fact that Peter has one child, Tony. In contrast, in an A-Box, the assertion only says that Tony is a child of Peter. Consequently, even if one also knows (by an assertion) that Tony is male, one cannot deduce that all of Peter's children are male. The only way of stating this is to make the assertion Peter$\leq$1 hasChild. This difference between databases and DL is important to remember when considering some of the later work.

2.8 OWL

OWL is the 'Web Ontology Language' and is a syntax for a variety of description logics. It is the focus of both academic and commercial work, and is the focus for most implemented DL reasoners. As can be seen from the diagram below, OWL builds on several layers of other technologies. However, I do not intend to cover all of these: OWL is a particular syntax for some description logics, but I concern myself solely with the logical aspects of OWL, rather than the syntax. Not only does this avoid us having to cover the whole set of technologies than underpin OWL, but it also spares us the OWL syntax. For those interested in the RDF/XML basis of OWL, see [58].

OWL comes in four 'flavours' - OWL-Lite, OWL-DL, OWL-Full and OWL-2.0. OWL-Lite is a fairly restricted subset of OWL (equivalent to $\mathcal{SHIF}(\mathcal{D})$), and as such is tractable and (relatively) easy to implement reasoners for. OWL-DL is equivalent to the description logic $\mathcal{SHOIN}(\mathcal{D})$, but remains decidable. OWL-Full is an extension of OWL equivalent to FOL. OWL-2.0 addresses some of the weaknesses in OWL-DL, while still remaining decideable, and is equivalent to the $\mathcal{SROIQ}(\mathcal{D})$ logic. Since it is much newer than OWL-Lite/ DL/ Full, some of the results mentioned below refer to OWL-DL rather than OWL-2.0. The relationship between OWL and some other recent developments in computer science and information technology is shown by the figure below. It shows how OWL uses more widely accepted technologies, such as XML, but is itself intended to serve as a foundation for other work in the future.

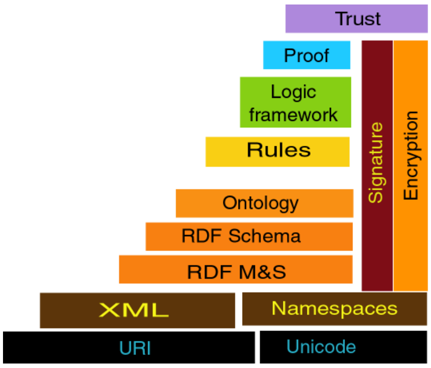

Figure 2.1: The Semantic Web Layer Cake. Note that the ontology layer builds on simpler formats, as as RDF and XML, and it is assumed that there will be other layers built 'on top' of the ontology.

2.8.1 Background

OWL-2.0 is an implementation of the $\mathcal{SROIQ}(\mathcal{D})$ description logic. $\mathcal{SROIQ}(\mathcal{D})$ is an expressive DL, and has a variety of allowed constructors that go beyond $\mathcal{AL}$. However, formulae in it can still be translated into FOL formulae with one free variable. I shall not give the details of the constructors allowed in $\mathcal{SROIQ}(\mathcal{D})$; we need only know that, like other DLs, it has a clearly defined syntax and semantics. For reasons that will become clear in the next chapter, we need the details only of ground literals. Before we can introduce these, however, I introduce one of the major differences between $\mathcal{AL}$ and logics such as $\mathcal{SROIQ}(\mathcal{D})$.

2.8.2 Datatypes

One of the requirements of many ontologies is to express information about real world datatypes, such as numbers, dates, times, etc. OWL allows us to do this by adding a second type of individual and role: datatypes and datatype roles. Datatypes are sets of typed numbers, as defined by the XML datatypes schemas. Datatype roles are binary relations between individuals and datatypes, which allow us to express facts such as a person's age (where the person would be an individual and their age some XML datatype). The details are not important here, but we will need to maintain the distinction between object roles (that relate concrete or abstract individuals) and datatype roles (which link concrete or abstract individuals and datatype individuals).

2.8.3 Ground Formulae in OWL

As we shall see later, it is the ground literals in an OWL ontology that shall interest us most. Because of this, I give an abbreviated definition of them here.

Definition 2.8.1. A $\mathcal{SROIQ}(\mathcal{D})$ vocabulary $\mathcal{K}$ consists of five sets of terms: C, D I, S, and R. In any vocabulary C and D are disjoint and R and S are disjoint, where the sets are:

C: The class names of a vocabulary, containing $\top$ and $\bot$

D: The datatype names of a vocabulary

R: The individual-valued property names of a vocabulary

S: The data-valued property names of a vocabulary

I: The individual names of a vocabulary

The definition of literals is similar in $\mathcal{SROIQ(D)}$ as to $\mathcal{AL}$, but includes datatype roles and negated object and datatype roles.

Definition 2.8.2. Let $\mathcal{K}$ be a vocabulary as above. Let $c \in C$ be a concept name, $r \in R$ an object property name, $s \in S$ a datatype property name, $d_1 \in D$ a datatype and $i_1, i_2 \in I$ concrete individuals. A $\mathcal{SROIQ(D)}$ atom is one of the following:

$c(i_1)$, $r(i_1, i_2)$, $s(i_1, d_1)$

A literal is an atom or a negated monadic atom, and the set of all literals in $\mathcal{SROIQ(D)}$ for a vocabulary $\mathcal{K}$ is denoted $\mathcal{K}^{\mathcal{A}}$.

Example 2.8.3. Let $\mathcal{K}$ be some vocabulary. Then if Men $\in C$, hasCat $\in R$, hasAge $\in S$, 40 $\in D$ Matt, Puss_in_Boots $\in I$, Men(Matt), hasAge(Matt,40) and hasCat(Matt, Puss_in_Boots) are atoms in $\mathcal{SROIQ(D)}$, and all three atoms, along with ¬Men(Matt), ¬hasAge(Matt,40), ¬hasCat(Matt, Puss_in_Boots) are literals. The set of all literals is denoted $\mathcal{K}^{\mathcal{A}}$.

Definition 2.8.4. Let $\mathcal{K}$ be some vocabulary. The set of non-ground formulae in $\mathcal{SROIQ(D)}$, denoted $\mathcal{K}^{\mathcal{T}}$, is the set of formulae given by the constructors defined by the letters of $\mathcal{SROIQ(D)}$.

Then we can define a $\mathcal{SROIQ(D)}$ ontology in a similar way as we defined an $\mathcal{AL}$ ontology.

Definition 2.8.5. Let $\mathcal{K}$ be a vocabulary. Then $\mathcal{O}_{\mathcal{K}} = \mathcal{K}^{\mathcal{A}} \cup \mathcal{K}^{\mathcal{T}}$ is the set of all formulae in $\mathcal{K}$. An ontology, $\Delta \subseteq \mathcal{O}_{\mathcal{K}}$ is some set of consistent, satisfiable formulae in $\mathcal{O}_{\mathcal{K}}$.

In the next section, it will be useful to be able to access the A-Box of an ontology directly, and we do this by a function.

Definition 2.8.6. The function $Literals_{\mathcal{K}}(\Delta)$ returns the set of literals in the A-Box of the ontology Δ.

2.8.4 DL ontologies and rules

$\mathcal{SROIQ(D)}$ is normally described as a very expressive DL. However, because of limits imposed due to a desire to maintain tractability, it cannot express certain concepts. The most important weakness is the lack of role constructors.

Example 2.8.7. Consider the roles parentOf and brotherOf. We might like to define some role, uncle, such that

$\forall x, y, z \; parentOf(x, y) \wedge brotherOf(y, z) \rightarrow uncleOf(x, z)$

Unfortunately, we do not have a role constructor in $\mathcal{SROIQ(D)}$, and so this is not permissible: inclusion of such constructors increases the complexity of the decision problem. In an effort to solve such problems there has been some work on the addition of rules to DL ontologies, but the nature of such rules, in terms of syntax and expressive power, remains a contested area.

2.8.4.1 SWRL

The most straight-forward approach is the W3C proposal for the 'Semantic Web Rule Language' or SWRL. SWRL is (again) a subset of FOL, and uses DL terms in the body and head of the rule, while restricting the head of the rule to single terms and assuming implicit universal quantification for the variables in the rule [40]. Unfortunately, SWRL is undecidable [39], and has also proved difficult to implement - the current approaches involved translating the ontology and rules to a set of FOL theorems and using a FOL theorem prover. As a response to this, there have been various approaches to defining subsets of SWRL that remain decidable. The one that concerns us here is DL-Safe rules.

2.8.4.2 DL-Safe Rules

[56] introduces a rule formalism that guarantees decidability. In the definitions given below, I follow the notation given in the paper.

Definition 2.8.8. Let $\mathcal{K}$ be some vocabulary and Δ some $\mathcal{SROIQ}(\mathcal{D})$ ontology. Then $\mathcal{O}(i)$ holds for any concrete individual $i \in I$.

$\mathcal{O}(i)$ is used to restrict the rules so that they apply to only the concrete individuals in the ontology.

Definition 2.8.9. Let A and B be single class or property names in the DL ontology or a non-DL atom. A rule r is DL Safe if it is of the form:

$$B_1 \wedge .... \wedge B_m \rightarrow A_1 \vee ... \vee A_n$$

and if each variable occurring in a DL-atom occurs in $\mathcal{O}()$ in the body of r, and each variable occurring in the head occurs in the body of the rule.

Example 2.8.10. Thus for the rules r_1 and r_2:

r_1 : Person(x) $\wedge$ livesAt(x,y) $\wedge$ worksAt(x,y) $\rightarrow$ Homeworker(x)

r_2 : $\mathcal{O}(x)$, $\mathcal{O}(y)$, Person(x) $\wedge$ livesAt(x,y) $\wedge$ worksAt(x,y) $\rightarrow$ Homeworker(x)

r_1 is not DL-safe, whereas r_2 is because of the non-DL atoms $\mathcal{O}(x)$ and $\mathcal{O}(y)$

The combination of the restrictions on DL-safe rules, both in terms of restricting variables in the head to appear in the body, and restricting the unification of the variables to the set of concrete individuals, provide a limited, but still useful rule language. The authors suggest that this is equivalent to appending the phrase "where the identity of all objects is known" to the normal reading of the rule, and they show that the combination of the $\mathcal{SHIQ}$ description logic and DL-safe rules is decidable, and make a case for believing that the same would be true with $\mathcal{SROIQ}(\mathcal{D})$. In practice, therefore, DL-Safe rules are suitable for reasoning over individuals in the ontology where such individuals are "known". However, we can see a shortcoming of such rules in this example:

Example 2.8.11. Let r_3 be a DL-safe rule:

r_3 : $\mathcal{O}(x)$, $\mathcal{O}(y)$, Married(x) $\wedge$ MarriedTo(x,y) $\rightarrow$ Happy(y) and let Married $\equiv$ Person $\exists$MarriedTo. Person be an axiom in the T-Box of an ontology, Δ_{Ex}. Then if the A-Box of Δ_{Ex} contains the following assertions:

Married(Matt), MarriedTo(Matt, Bhanu)

we can infer

Happy(Bhanu)

However, if the A-Box of Δ_{Ex} only contains the assertion

Married(Matt)

then we can infer (from the ontology alone) that there exists a person to whom Matt is married, and so

MarriedTo(Matt, i_1)

holds, but since their identity is not known, i_1 is not a concrete individual, and so we cannot apply the DL-safe rule to them.

2.8.4.3 Rules in Argumentation

One of the core features of argumentation is the construction of arguments. In the next chapter, we shall use defeasible rules for the construction of arguments, and towards the end of that chapter we shall address to what extent my defeasible rules are similar to DL-Safe rules.

2.8.5 Tools

One of the motivations behind the development of OWL was to enable the production of toolsets to help in the production of sizable ontologies. These are mainly divided into author/viewing tools, and reasoners. There are many tools in both categories, and I shall briefly mention one from each, namely Protege [48]and Pellet [70].

Protege is developed and maintained by the Stanford Medical Informatics group. Originally developed as an editor for Frame-Based systems, much of its use now revolves around OWL. It provides a graphical interface for defining concepts and roles, as well as the ability to enter logical definitions for classes as text. Notable recent additions have included the addition of a ‘controlled natural language’ style syntax for class definitions and query facilities.

Pellet implements the tableau-reasoning algorithm for OWL-2.0 as described in [36]. It is currently the only freely available reasoner that fully implements OWL-2.0 and implements DL-Safe rules. Although the details of these implementations are not important, it is useful to be aware that implementations exist, and that we can use them in the work. For example, although I give a graphical representation of the ontological class structure below, and an axiomatic representation in appendix A, the best way to explore the ontology is by downloading a copy of the OWL file and viewing it in Protege.

2.9 A Breast Cancer Ontology

2.9.1 Introduction

My aim in developing the Breast Cancer ontology is not to provide a perfect, generic ontology which encompasses all of breast cancer. Instead, I have based it on relevant literature, fleshed out on the basis of personal experience. For example, there is little mention in the literature of ‘treatment’ in general - it refers to specific treatments. However, to include the specific treatment, but omit the more general term would seem perverse, and so I have added it.

2.9.1.1 Information Sources

I started by constructing a first draft of the ontology on the basis of experience, and then refined it based on terms drawn from the papers used in the worked example in chapter 8; for full details on paper selection, please see there. In brief, they are a set of references from the the US National Cancer Institute (NCI), restricted to those dealing with Tamoxifen patients with invasive breast cancer.

2.9.2 The Ontology

Presenting an ontology rich enough to inform the work is difficult; The best way to understand it is to view it in a tool such as Protege. It contains 199 classes, of which 12 are defined, and 26 properties, of which 23 are object properties. For reasons of readability, I give here only the class hierarchy (in graphical form), but the ontology is best understood by downloading it and using a tool such as Protege to explore it.

The class structure of the ontology is shown below. To make it manageable, I present an overview of the entire ontology, and then seven partitions in more detail Fig:2.9. Each diagram should be read left to right (superclass to subclass); the boxes around each section of the larger ontology show which smaller box expands that section, and lighter coloured classes are undefined; darker ones have an axiomatic definition in the ontology. Classes that have a small black triangle at the right-hand end of the ellipse have subclasses that have been hidden in that diagram, those with a triangle at the left hand end have hidden superclasses. I start by presenting a figure that although it has details too small to read, shows the general shape and scale of the ontology. Details are given in the following pages. In this and following chapters, the breast cancer ontology will be referred to as Δ_{BC}.

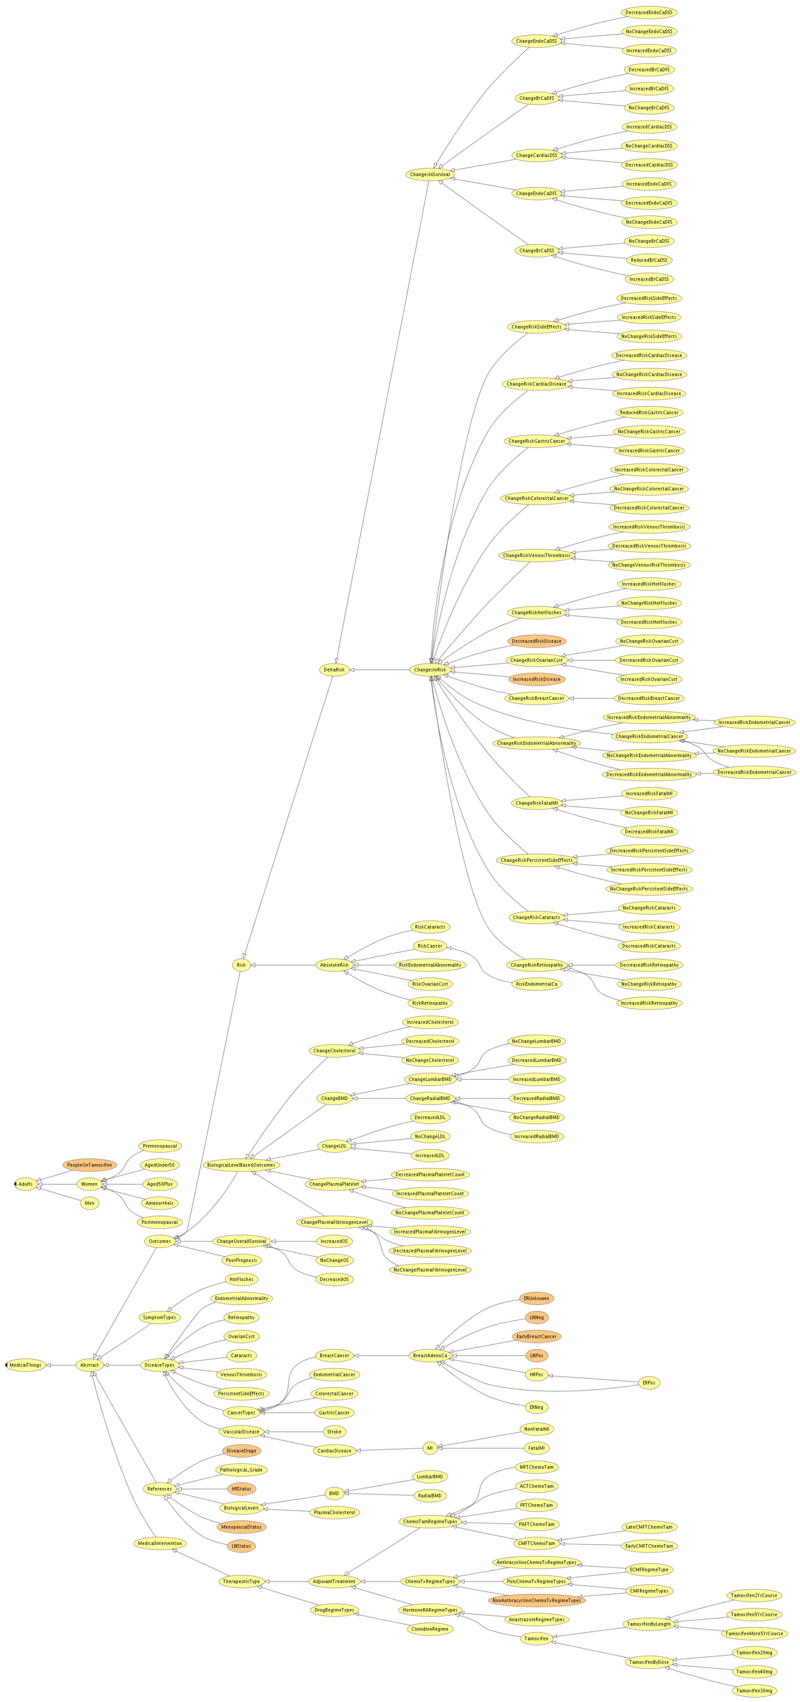

Figure 2.2: The Breast Cancer Ontology. Although impossible to read at this scale, the diagram shows the ontology as a whole. Different areas of the ontology are expanded below.

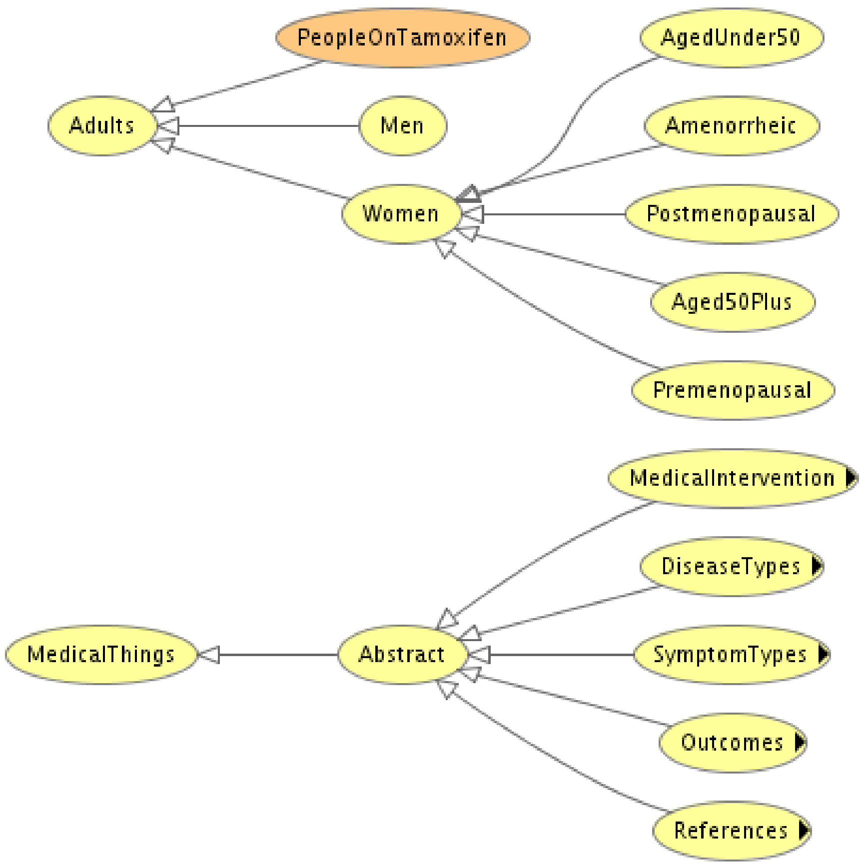

Figure 2.3: Upper three levels of the ontology. The black triangles at the right-hand end of a class indicate subclasses, hidden in this image

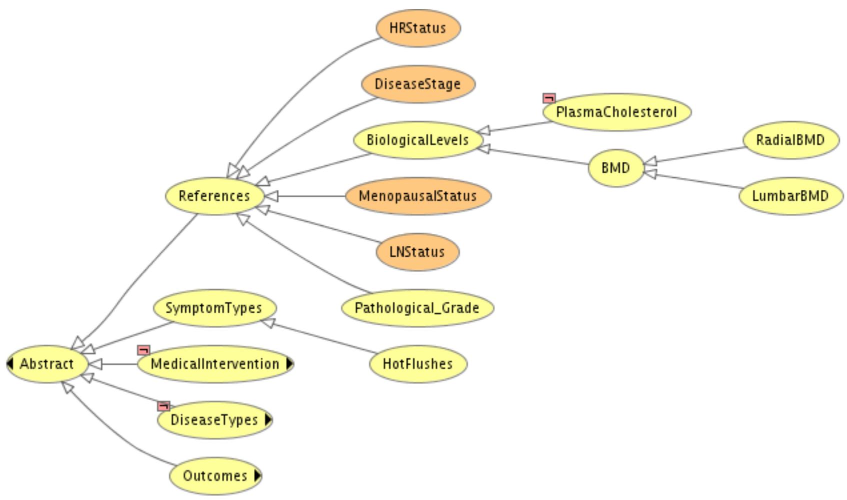

Figure 2.4: Expanded References and Symptom Types

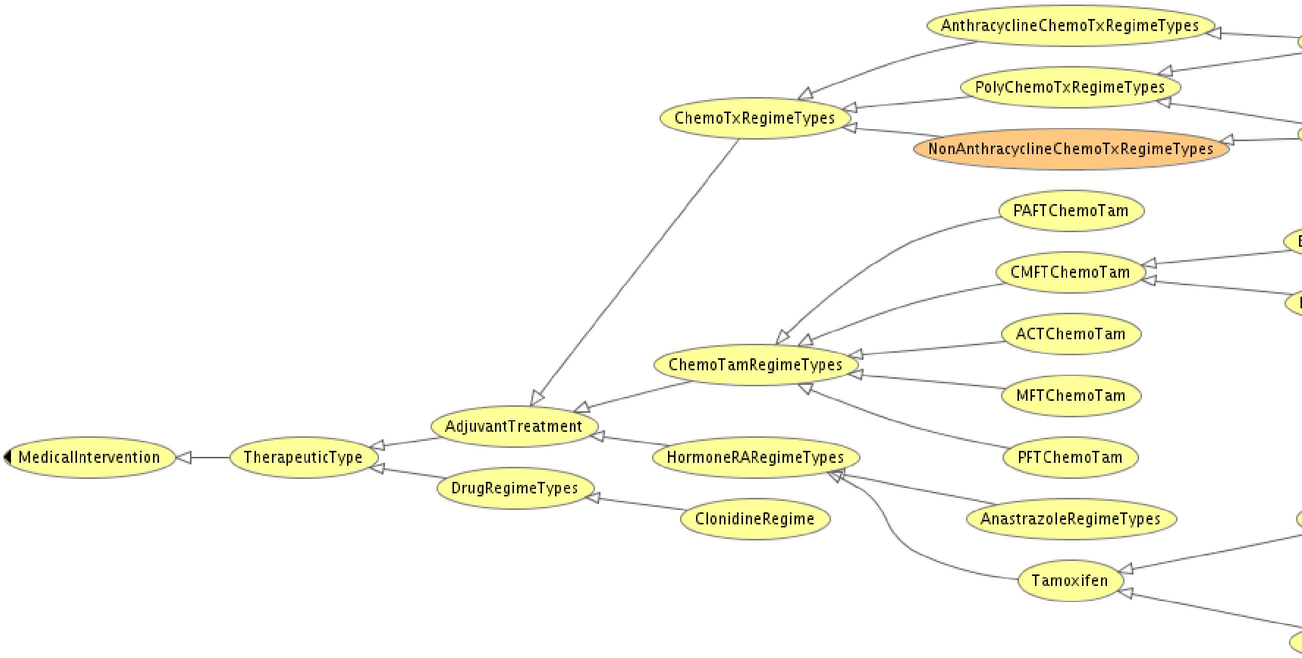

Figure 2.5: The subclasses of Medical Interventions. Classes in dark orange are those that are have logical definitions; most classes are defined only in relation to their sub/super-class relationships.

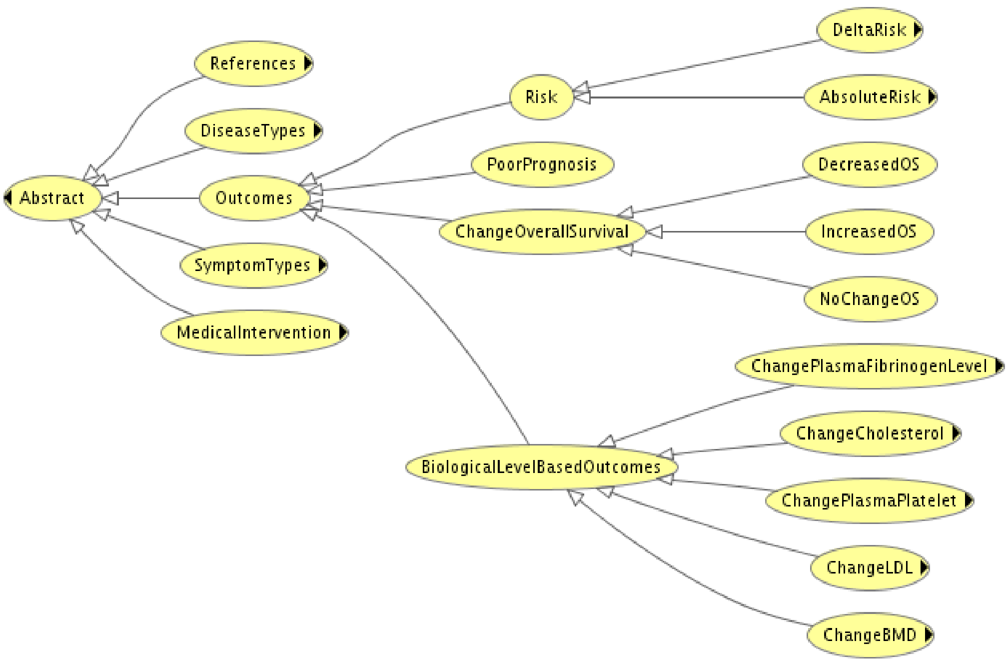

Figure 2.6: The subclasses of Outcome and Risk

Figure 2.7: The subclasses of Biological levels based outcome

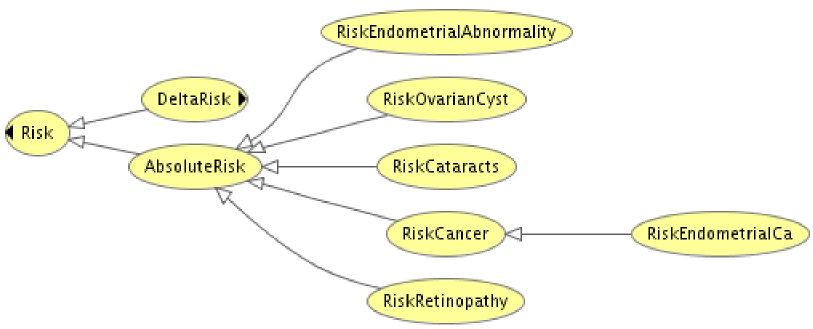

Figure 2.8: The subclasses of Absolute Risk. The black triangle to the left of a class name indicates hidden super classes; to the right of a class name indicates hidden subclasses.

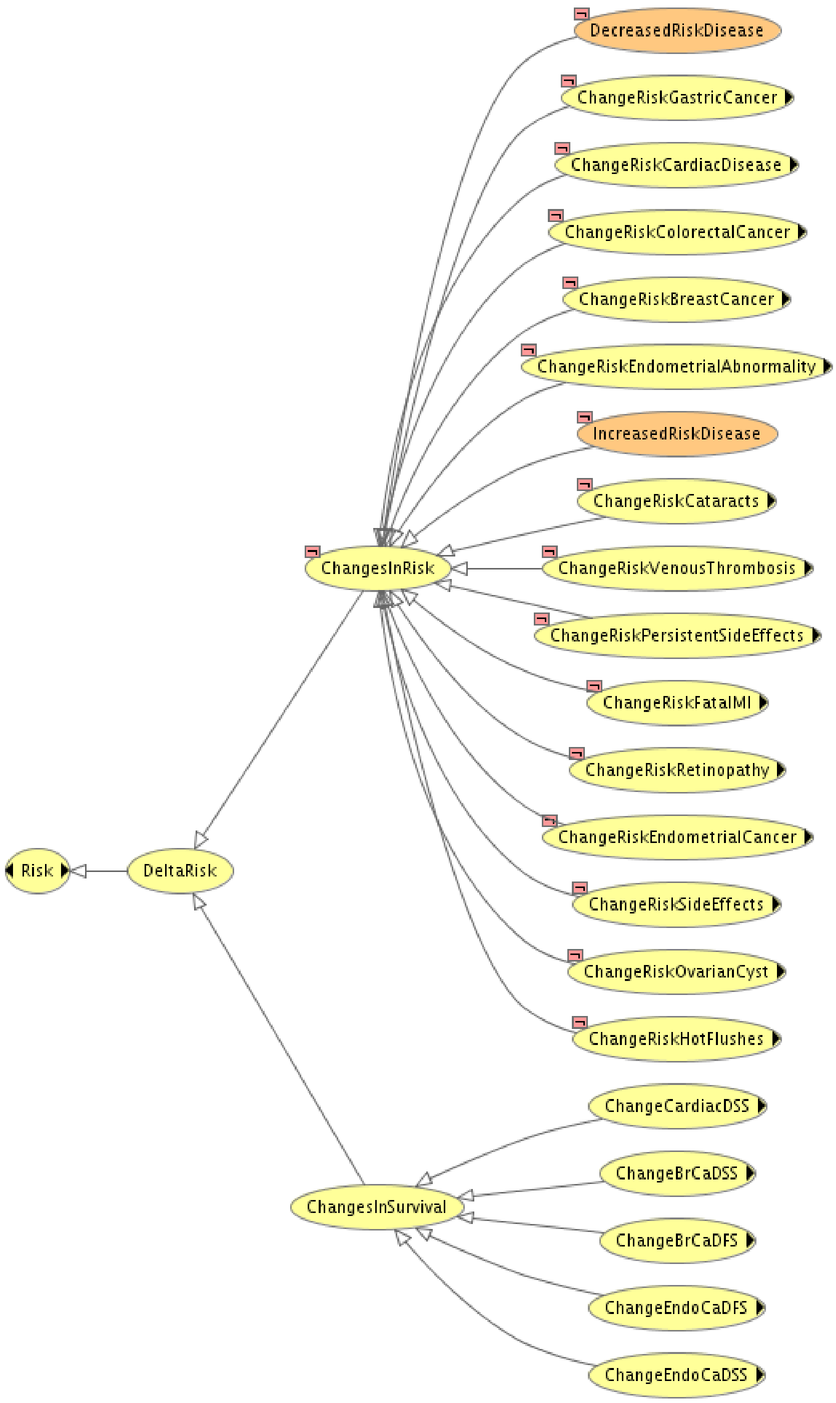

Figure 2.9: The classes of Changes in Risk and Changes in Survival. Each specific change in risk or survival has three subclasses, representing an increase, decrease or no change in that outcome.

2.9.3 Some modelling choices

As with any knowledge representation formalism, I have had to make a set of modelling decisions. Some of these may seem intuitively ‘correct’; some may seem arbitrary, but I explain the rationale for these decisions below.

2.9.3.1 Designing an Ontology for a Reason

If we were designing an ontology for an electronic patient record, we would need a much more complex structure. Such an ontology is unlikely to be necessary for my purposes, and of course involves the additional work (on both my and the reader's part) with little benefit for the thesis. Instead, my ontology has people, disease types, treatment types and a few references. Such simplification is not essential, but allows me to present the work in a slightly more compact fashion; again, I hope to demonstrate that the underlying idea is sound, not to produce a system that is ready to support an NHS clinic, and here a little brevity will hopefully lead to clarity. The final omission is of precise doses - where possible, I have referred to 'standard courses'. In some cases these are defined, but in other the associated mathematics (for example working out the dose of a chemotherapy regime given a patient's surface area) is of little interest to us here, and so is omitted.

2.9.3.2 Individuals as References

In the ontology, I have used different individuals as references. It is easy to associate concrete individuals in an ontology with concrete objects in the world - individual people or cars, for example. However, they are also used to model unique occurrences of abstract contexts. Consider the colour red in some colour ontology. When two people talk of the colour red, they need to talk of the *same* colour. Here we have concrete (i.e. named) individuals, referring to abstract ideas in order to ensure referential consistency. In the ontology, we have similar abstract, atemporal concepts, that are unique. That is, when I talk of a tumour that has TNM tumour stage of `T1`, it is the same `T1` that other tumours have, even though the tumours may be different, and even though, for different tumour types, they may be of radically different sizes (e.g. For adenocarcinoma of the breast, tumours of <2cm are T1; for glottic laryngeal tumours, T1 tumours are those involving one or both cords, as long as both cords remain mobile - the definitions are very different, but they both have a T stage of T1). This principle, of using individuals as references, is then further extended for treatments and diseases.

2.9.3.3 From Drugs to Treatments

While there might be clear ontological representations of (let us say) a class of drug, as a clinician I do not simply prescribe someone an instance of this class. Instead, I prescribe some treatment course, which has a specified dose, frequency and duration.

Example 2.9.1. When I see someone in clinic and prescribe them Tamoxifen, I do not just say "Have some Tamoxifen". I prescribe 20mg Tamoxifen, to be taken orally once per day for five years.

This sort of information is even more important in the use of complex treatments such as Chemo- and Radio-therapy. For example, breast cancer patients may be prescribed ECMF chemotherapy, which is a combination of Epirubicin, Cyclophosphamide, Methotrexate and 5-Fluorouracil (hence the abbreviation). But these drugs are not given once: Instead the standard course would be to give Epirubicin once every 3 weeks for 4 cycles, followed by 4 cycles of combined CMF. The dose of each drug is normally calculated for each patient, based either on weight or surface area. Now, we might easily have an ontological representation of the individual drugs, but to give a representation of a course of ECMF treatment is more challenging. For clinicians, combination drugs regimes (such as ECMF) are the basic ‘atoms’ with which we build treatment courses, and debates about treatments revolve around which treatment courses to use. However, each patient is obviously prescribed a different instantiation of any treatment course, and this has the potential to lead to a profusion of near-identical courses of treatment. Such detail would be essential if we wanted to build a real-world electronic prescribing system. Since we do not, I again use reference individuals to describe both treatment courses and diseases.

2.9.3.4 Prototypes

Example 2.9.2. Consider the Tamoxifen example above; I might prescribe many courses of “Tamoxifen 20mg 1/day for 5 years” over a year, but they all relate to the same ideal treatment. The instantiations of these might be made by different manufacturers, at different times, and given to different people, but they all refer to the same (unique) ideal.

Arguing about these ‘idealised’ courses is both more parsimonious and also fits better with the literature (they similarly report the results of typical couses of treatment). I refer to the idea of 2 or 5 years worth of Tamoxifen at 20mg per day as a ‘prototype’ : it represents a common aspect of all the instances of the treatment that are taken by patients. These prototypes can themselves be organised in an ontological hierarchy:

Example 2.9.3. We could consider two different types of tamoxifen treatment, two and five years worth of treatment. We might then have two prototypes, TamD2Yr and TamE5Yr. However, both would be instances of the class of TamoxifenTreatments.

There is a similar argument to be made about the representation of diseases. Actual instances of disease would clearly need a great deal of information to fully describe them. However, we will often want to refer to disease prototypes, in the same way as we refer to treatment prototypes, and similarly we can arrange such instances in an ontological hierarchy.

2.9.3.5 Rules

In the next chapter, we will see how conflict between arguments is defined as logical inconsistency with respect to the ontology. However, in order to ensure that we can do this effectively in a logic such as $\mathcal{SROIQ}(\mathcal{D})$, I need to step outside DL inference and augment the ontology with DL-Safe rules. I shall therefore refer to the ontology to mean the complete set of strict inferences, those based on both DL and rules. Such compromises seem inevitable given my desire to incorporate both ontology and argumentation in one system, and represent an interesting comment on the expressiveness of DL ontologies. For the work in the thesis, the ontology presented above is augumented with rules of the form shown below, where predicates are concepts and roles in the ontology and $\perp$ denotes contradiction:

$$People(x) \wedge hasDeltaRisk(x,y) \wedge IncreasedOS(y) \wedge hasDeltaRisk(x,z) \wedge DecreasedOS(z) \Rightarrow \perp$$

$$People(x) \wedge hasDeltaRisk(x,y) \wedge IncreasedDFS(y) \wedge hasDeltaRisk(x,z) \wedge DecreasedDFS(z) \Rightarrow \perp$$

For all of the pairwise disjoint changes in risk (increased OS/ decreased OS; increased DFS/ decreased DFS; increased risk cardiac disease/ decreased risk of cardiac disease; etc.). These rules are necessary because of a requirement in the way we want to reason about risk. Although it is simple to specify that Increased and Decreased OS are disjoint classes, we need to impose a limitation that one person should not have both an increase and decrease in either OS or DFS; one solution would appear to be to ensure that each person is only related to one change in risk via hasDeltaRisk (i.e. make hasDeltaRisk a functional property); however, one person can have an expectation of both increased DFS and decreased OS, and so that would be too restrictive. One solution is to use the type of rules described above.

2.9.3.6 Encoding Intention

In chapter 5, I shall deal with intention. The approach laid out in this thesis is designed to work with a variety of different ontologies in different domains, and so I have tried to take a relatively liberal approach in specifying what must be present or not present in the ontology. However, since one of the things I shall specify in later chapters is that all the terms that we argue about must be present in the ontology, for us to be able to argue about intention, there must be some concepts and roles in the ontology that reflect ideas of intention, and these are:

Definition 2.9.4. For some DL ontology Δ, Δ must contain the role hasPosIntent and hasNegIntent where:

> The range of hasPosIntent/ hasNegIntent is People, and the domain of hasPosIntent/hasNegIntent is Treatments. hasNegIntent and hasPosIntent are both functional properties, and given two individuals $i_1, i_2 \in I$, {hasNegIntent(i_1, i_2) $\cup$ hasPosIntent(i_1, i_2)} $\vdash_{Ont} \perp$

2.10 Summary

I have briefly reviewed the history, underlying logic and syntax of OWL, the W3C standard for ontologies, as well as presenting my breast cancer ontology and detailing some of the choices behind the modelling choices made during its construction. I now move onto the presentation of the argumentation system.

Chapter 3

Ontology-based Argumentation Framework

I present my proposal for an ontology-based argumentation framework, (OAF), and descibe some of its properties. I then go on to consider a small example that illustrates some of its features. OAF is based heavily on DeLP [32], but, unlike DeLP, all the predicate and constant symbols used must be present in the ontology, and unlike DeLP, there are no strict rules, as non-defeasible inferences are captured by ontological inference. I start by giving definitions of a simple logical language, based on the elements of the ontology, and use these to define defeasible rules. I then introduce definitions of defeasible and ontological reasoning, and demonstrate how both are used to develop arguments, and conflict between arguments is defined in terms of unsatisfiable conjunctions of formulae with respect to the ontology. This chapter concentrates on the core definitions of OAF, while the next two chapters define further extensions to the system. Therefore, the version of OAF presented in this chapter is referred to as 'simple' OAF (S-OAF).

3.1 The Language

An OAF is a set of defeasible rules, and a DL ontology, such that the atoms used in the rules are found in the ontology. The ontology provides a 'model' of the world, both by defining terms in relation to one another (the T-Box) and providing information about individuals (the A-Box), while rules make defeasible assertions about individuals in the world from which we construct arguments.

My requirements are relatively simple; I am content to argue about ground instances (and so do not need full FOL) but at the same time the rules need to be written in a form that allows for multiple different groundings - I do not want to have to write rules about every person in the ontology. At the same time, I want the rules to be 'compatible' with the ontology, in that they need to use the same predicates and constants, so that we can use the ontology both as a source of information about individuals and as an 'oracle' to answer questions about entailment and conflict. Finally, I am interested in only a simple form of entailment in the rule language (essentially a form of modus ponens).

In order to do this, therefore, I start by defining a language that is a propositional logic with variables, and relate this to some ontological language $\mathcal{K}$. In my case, this will be $\mathcal{SROIQ}(D)$, but the definitions in OAF do not prescribe this.

Definition 3.1.1. Let $\mathcal{K}$ be a vocabulary, as presented in the last chapter such that for some DL ontology language, $\mathcal{K}$ as presented in 2.7.1:

- C be the set of concept names in $\mathcal{K}$
- I be the set of concrete individuals in $\mathcal{K}$, denoted a, b, c, possibly with subscripts
- D be the set of datatype value names in $\mathcal{K}$ denoted $d_1...d_n$
- R be the set of individual-valued role names in $\mathcal{K}$
- S be the set of datatype rolenames in $\mathcal{K}$

In addition, the vocabulary contains the logical symbols $\neg$ and $\wedge$. Note that in this and subsequent chapters, I use $\wedge$ instead of the $\sqcap$ symbol that is normally used for description logic.

Given these definitions for the language, we can now define our atoms, literals and formulae.

Definition 3.1.2. A ground ***atom*** is any of the following:

- $p(t)$where $p \in C$ and $t \in I$
- $p(t_1, t_2)$ where $p \in R$ and $t_1, t_2 \in I$
- $p(t_1, t_2)$ where $p \in S$ and $t_1 \in I$, $t_2 \in D$

Example 3.1.3. Let $\mathcal{K}$ be a vocabulary such that C is the set of concept names in $\mathcal{K}$ and I is the set of individuals in $\mathcal{K}$ such that Men $\in C$ and Matt $\in I$. Then Men(Matt) is an atom.

Definition 3.1.4. A ***literal*** is a ground atom or a negated ground atom (i.e. if $p(t)$ is an atom, then $p(t)$ and $\neg p(t)$ are literals). The set of all literals is denoted $\mathcal{G}_{\mathcal{K}}$.

Note that the definition of literals matches that of literals in OWL, and so literals in our language are in fact the ground formulae ($\mathcal{A}^{\mathcal{I}}$) in our DL. This is why I allow literals of the form $S_1(a, d_1)$ but not of the form $S_1(d_1, a)$, as this is not a literal in OWL. Note also that all atoms (and hence literals) are ground. I now give a definition of formulae.

Definition 3.1.5. For some vocabulary $\mathcal{K}$, the set of ***conjunctive formulae*** is defined such that any literal is a formula, and if α and β are formulae, then $\alpha \wedge \beta$ is a formula. The set of all formulae is denoted $\mathcal{F}_{\mathcal{K}}$.

Example 3.1.6. For some concept names C_1... C_n, individual-valued role names R_1.... R_n, datatype rolenames S_1... S_n, individuals a....c and datatypes d_1.... d_n the following are formulae in $\mathcal{F}_{\mathcal{K}}$:

$C_1(a) \wedge C_2(a)$

$C_1(a) \wedge\neg C_3(b)$

$C_1(a) \wedge R_1(a, b) \wedge C_2(b)$

$C_1(a) \wedge R_1(a, b) \wedge S_1(a, d_1)$

Note that in the last chapter, I introduced the set of formulae with reference to some vocabulary $\mathcal{K}$ and the DL, $\mathcal{SROIQ}(\mathcal{D})$, denoted $\mathcal{O}_{\mathcal{K}}$. Here I introduce a different set of formulae, $\mathcal{F}_{\mathcal{K}}$, and although the sets of ground literals in both languages are the same, the sets of formulae are not the same - that is $\mathcal{O}_{\mathcal{K}} \neq \mathcal{F}_{\mathcal{K}}$.

I now define some functions that will be of use in later work.

Definition 3.1.7. For the set of formulae, $\mathcal{F}_{\mathcal{K}}$ and literals $\mathcal{G}_{\mathcal{K}}$, the function $Atoms(\phi) : \mathcal{F}_K \mapsto \mathcal{G}_{\mathcal{K}}$ takes a formula, ϕ and returns all the atoms occurring in the subformulae of ϕ.

Example 3.1.8. Consider the formulae:

ϕ_1 = People(MsJones) $\wedge$ $\neg$Happy(MsJones).

Then $Atoms(\phi_1)$ = {People(MsJones),Happy(MsJones)}

ϕ_2 = People(MsJones) $\wedge$ hasDisease(MsJones,ProtoBreastCancer) $\wedge$ BreastCancer(ProtoBreastCancer).

Then $Atoms(\phi_2)$ = {People(MsJones), hasDisease(MsJones,ProtoBreastCancer), BreastCancer(ProtoBreastCancer)}

I now define some other useful functions.

Definition 3.1.9. For some set of formulae, $\mathcal{F}_{\mathcal{K}}$, the function $Classes(\phi)$: $\mathcal{F}_{\mathcal{K}} \mapsto \mathcal{G}_{\mathcal{K}}$ takes a formula and returns all the monadic atoms occurring in the subformulae of ϕ.

Example 3.1.10. Consider the formula:

ϕ_1 = People(MsJones) $\wedge$ hasDisease(MsJones,ProtoBreastCancer) $\wedge$ BreastCancer-Types(ProtoBreastCancer).

Then $Classes(\phi_1)$ = {People(MsJones),BreastCancerTypes(ProtoBreastCancer)}

Definition 3.1.11. The function $Literals(\phi)$: $\mathcal{F}_{\mathcal{K}} \mapsto \mathcal{G}_{\mathcal{K}}$ takes a formula and returns all the literals occurring in the subformulae of ϕ.

Example 3.1.12. Consider the formula ϕ_1 = People(MsJones) $\wedge$ $\neg$Happy(MsJones).

Then $Literals(\phi_1)$ = {People(MsJones),$\neg$Happy(MsJones)}

Note that this is distinct from the function $Literals_{\mathcal{K}}(\Delta)$ given in the last chapter 2.8.6, although it is related: if $\phi = \bigwedge \Delta^{\mathcal{A}}$, then $Literals(\phi_1) = Literals_{\mathcal{K}}(\Delta)$.

Definition 3.1.13. The function $Constants(\phi)$: $\mathcal{F}_{\mathcal{K}} \mapsto I \cup D$ takes a formula and returns the individuals and datatypes occurring in subformulae of ϕ.

Example 3.1.14. Consider the formula ϕ_1 = People(MsJones) $\wedge$ hasAge(MsJones,40) $\wedge$ hasDisease(MsJones,ProtoBreastCancer). Then $Constants(\phi_1)$ = {MsJones,40,ProtoBreastCancer}

Now we have these basic definitions, I can define how we construct rules. A rule consists of two formulae, connected via defeasible implication, where the body of the rule is either a formula or null, and the head of a rule is a literal. The definition is as follows:

Definition 3.1.15. For some vocabulary $\mathcal{K}$, let $\mathcal{F}_{\mathcal{K}}$ be the set of conjunctive formulae, $\mathcal{G}_{\mathcal{K}}$ be the set of literals in $\mathcal{K}$ and $\mathcal{B}$ be a set of alphanumeric labels.

If $\alpha \in \mathcal{G}_{\mathcal{K}}$, $L \in \mathcal{B}$ then $L : \Rightarrow \alpha$ is a defeasible rule

If $\alpha \in \mathcal{G}_{\mathcal{K}}$, $\beta \in \mathcal{F}_{\mathcal{K}}$, $L \in \mathcal{B}$ and $Constants(\alpha) \subseteq Constants(\beta)$ then L: $\beta \Rightarrow \alpha$ is a defeasible rule

The set of all defeasible rules is denoted $\mathcal{R}_{\mathcal{K}}$.

I refer to β as the ***Body*** of the rule, α as the ***Head*** of the rule and L as the ***Label.*** The label refers to the source of the information that was used to create the rule and it will sometimes be useful to refer to a rule by its label, rather than restating the rule. The constants in the head of a rule are a subset of those in the body in order to help with type checking of individuals and to allow for the development of schematic rules (see below).

Example 3.1.16. The following is an example of a defeasible rule concerned with the effect of having breast cancer on overall survival (OS):

NEJM2001:People(MsJones)

$\wedge$ hasDisease(MsJones, ProtoBreastCancer)

$\wedge$ DecreasedOS(DecreasedOS0.7)

$\Rightarrow$ hasDeltaRisk(MsJones,DecreasedOS0.7)

It is worth making three remarks here: firstly, sans serif font (as in People(MsJones)) is used to denote ontological terms. Secondly, given the length of my rules, where necessary I shall tend to break the line at either $\wedge$ or $\Rightarrow$ symbols to increase readability. Thirdly, although there is no restriction on the format of a rule's label, I will tend to use labels that have some meaning - in the majority of the thesis, this means that they refer to some published study, and here I am assuming that the rule above has been developed from some study published in the New England Journal of Medicine (NEJM) in 2001. Chapter 6 describes a simple technique for using these labels to access information about the study from which the rule was drawn, and hence decide on preference status between rules and arguments.

For reasons of brevity, it will sometimes be useful to refer to a rule by its label, rather than restating the whole rule. It will also be useful to be able to refer to the different parts of the rule through some function. I start by describing some functions to allow us to check the literals that make up the rules.

Definition 3.1.17. The $Head$ and $Body$ functions are defined as follows: for some rule $r \in \mathcal{R}_{\mathcal{K}}$, where r is a defeasible rule, if r is of the form: $L_1 : \Rightarrow \alpha$, then $Body(r)$ = null and $Head(r)$ = α. If r is of the form: L_2: $\beta \Rightarrow \alpha$, then $Body(r)$ = β and $Head(r)$ = α.

Example 3.1.18. Given a rule of the form:

r_1: People(MsJones)

$\wedge$ hasDisease(MsJones, ProtoBreastCancer)

$\wedge$ BreastCancerTypes(ProtoBreastCancer)

$\Rightarrow$ IncreasedMortality(MsJones)

$Body(r_1)$= People(MsJones) $\wedge$ hasDisease(MsJones,ProtoBreastCancer)

$\wedge$BreastCancer(ProtoBreastCancer)

$Head(r_1)$ = IncreasedMortality(MsJones)

3.2 Representing Knowledge

I have introduced the language of OAF above. This section defines how I use this language to capture our domain knowledge. I start by recalling the definition of an ontology.

Definition 3.2.1. For some vocabulary $\mathcal{K}$, the set of all formulae in $\mathcal{K}$ is denoted $\mathcal{O}_{\mathcal{K}}$. An ***ontology*** $\Delta \subseteq \mathcal{O}_{\mathcal{K}}$ is a consistent, satisfiable set of well-formed formulae (according to some DL) in $\mathcal{K}$.

Given some ontology and an associated set of rules, we can define an ontology-based argumentation framework.

Definition 3.2.2. Let $\mathcal{K}$ be some vocabulary, $\mathcal{O}_{\mathcal{K}}$ be the set of ontological formulae in $\mathcal{K}$ and $\mathcal{R}_{\mathcal{K}}$ be the associated set of defeasible rules. A ***Simple Ontology-based Argumentation Framework*** (S-OAF), denoted Ω, is a pair (Θ,Δ) where $\Delta \subseteq \mathcal{O}_{\mathcal{K}}$ is a DL ontology and $\Theta \subseteq \mathcal{R}_{\mathcal{K}}$ is a set of defeasible rules, referred to as a rulebase.

The examples in the thesis all relate to the breast cancer ontology, which I name here.

Definition 3.2.3. The breast cancer ontology (presented in the last chapter) is denoted Δ_{BC}.

Defeasible rules are ground. However, following the convention used by other authors [32, 51], some examples will use 'schematic rules' with variables. Given some 'schematic rule' r, the schematic rule stands for the set of all ground instances of r where, for some DL ontology language $\mathcal{K}$, the constant symbols are from $I \subset \mathcal{K}$. In order to distinguish variables in schematic rules, they will be denoted with $w..\ z$, possibly with subscripts. I illustrate the use of schematic rules with an example.

Example 3.2.4. Let $\mathcal{K}$ be some vocabulary and $\mathcal{G}_{\mathcal{K}}$ the set of literals. Let $\Gamma \subset \mathcal{G}_{\mathcal{K}}$ be a set of literals and recall that Δ_{BC} is the breast cancer ontology (as defined in the last chapter). Let Ω = (Θ, $\Delta_{BC} \cup \Gamma$) be a S-OAF where Θ and Γ are defined as:

Γ= {People(MsSmith), hasDisease(MsSmith, ProtoBreastCancer), People(MsJones), hasDisease(MsJones, ProtoBreastCancer)}

Θ = $\{r_1\}$

where r_1 is the schematic rule:

r_1: People(x)

$\wedge$ hasDisease(x,ProtoBreastCancer)

$\wedge$ BreastCancerTypes(ProtoBreastCancer)

$\Rightarrow$ IncreasedMortality(x)

Then given Γ and Δ_{BC}, the above schematic rule stands for:

r_1^1: People(MsJones)

$\wedge$ hasDisease(MsJones, ProtoBreastCancer)

$\wedge$ BreastCancerTypes(ProtoBreastCancer)

$\Rightarrow$ IncreasedMortality(MsJones)

r_1^2: People(MsSmith)

$\wedge$ hasDisease(MsSmith, ProtoBreastCancer)

$\wedge$ BreastCancerTypes(ProtoBreastCancer)

$\Rightarrow$ IncreasedMortality(MsSmith)

The use of schematic rules allows for rules that apply to all the members of some group (in this case, people). This offers some economies of presentation and implementation.

The work presented in this thesis is mainly concerned with rules that concern risk (epistemic rules), rules about intention (ergonic rules) and rules about treatments (heuristic rules). The difference between these rules lies in their use of different atoms in the head and body of the rules, and they are used to represent different types of knowledge.

Definition 3.2.5. Let $\mathcal{K}$ be some vocabulary and I be the set of named individuals in $\mathcal{K}$. For some $t_1, t_2, t_3 \in I$, a defeasible rule, r, is an ***epistemic*** defeasible rule iff the two following conditions hold:

1. $Head(r)$ = hasDeltaRisk(t_1,t_2) or $Head(r)$ = hasRisk(t_1,t_2)

2. There is at most one atom of the form hasTreatment(t_1,t_3) $\in Body(r)$

The set of all epistemic rules is denoted $\mathcal{R}_{\mathcal{K}}^{\mathcal{P}}$.

Therefore, an 'epistemic' rule is not just one about belief; it uses a particular selection of atoms from the language. Specifically, epistemic rules are about risk - either absolute risk levels, or changes in risk - possibly in combination with treatments.

Definition 3.2.6. Let $\mathcal{K}$ be some vocabulary and I be the set of named individuals in $\mathcal{K}$. For some $t_1, t_2 \in I$, a defeasible rule, r, is an ***ergonic*** defeasible rule iff: $Head(r)$ = hasPosIntent(t_1,t_2) or $Head(r)$ = hasNegIntent(t_1,t_2). The set of all ergonic rules is denoted $\mathcal{R}_{\mathcal{K}}^{\mathcal{R}}$.

I shall not discuss ergonic rules any further here; they are considered in detail in the next chapter.

Definition 3.2.7. Let $\mathcal{K}$ be some vocabulary and I be the set of named individuals in $\mathcal{K}$. For some $t_1, t_2 \in I$, a defeasible rule, r, is a ***heuristic*** defeasible rule iff: $Head(r)$ = hasTreatment(t_1,t_2) or $Head(r)$ = $\neg$hasTreatment(t_1,t_2). The set of all heuristic rules is denoted $\mathcal{R}_{\mathcal{K}}^{\mathcal{H}}$.

The uses of heuristic rules will become illustrated later in this chapter. Although such rules may seem natural, there are problems associated with their use. Not all rules in $\mathcal{R}_{\mathcal{K}}$ are heuristic, epistemic or ergonic, but these three sets are pairwise disjoint.

3.3 Ontological Reasoning in OAF

At the end of the last chapter, I said that we would deal with the relationship between the ontology and the argumentation system in this chapter, and this is what I now turn to. So far, I have used the concept, role, datatype value and individual names in the system, and one of the aims of this approach is to allow us to proceed with our argumentation without worrying about the precise details of the ontological formalism. While I assume some basic properties (for example, that our ontology is DL based, and that it contains notions of conflict, entailment, etc.). I do not specify any particular algorithm or approach to satisfying these requirements, thereby allowing us to take advantage of advances in work on DL ontologies without changing our approach. This approach also fits with the idea that I regard myself as being a ‘consumer’ of ontologies and ontology research, rather than directly contributing to the field. In particular, any weaknesses and problems with the ontological aspects of the work will generally be either left as a topic for future work or worked around. Furthermore, trying to handle both ontologies and argumentation in the same formalism makes it far harder to compare this work on argumentation to other works, and so risks disguising the contribution to argumentation that this thesis makes. For all these reasons, I have sought to encapsulate the ontology. This approach is of course similar to the idea of writing software “to an interface”. Before we can proceed to this, though, I need to handle the issue of translation.

3.3.1 From DL to OAF

From the perspective of OAF, the ontology is used for two purposes. First, OAF uses the ontology as a dictionary and syntax: it contains the names that I may use to form atoms, as well as imposing some simple constraints on their syntax. So the ontology tells us in broad terms what predicates we can use in our rules, but also contains specific information (in the form of literals) about certain instances.

Example 3.3.1. Consider the DL ontology Δ_{BC}, where the A-Box contains the following assertion:

Tamoxifen5YrCourse(TamE5Yr)

Then this is also a literal in $\mathcal{G}_{\mathcal{K}}$.

The second use of the ontology in Δ is to provide an entailment relation.

Definition 3.3.2. The symbol $\vdash_{Ont}$ is used to denote ontological entailment in $\mathcal{K}$ as defined in 2.7.5.2 and 2.7.5, so $\Delta \vdash_{Ont} \phi$ denotes that ϕ is entailed by the ontology Δ.

Example 3.3.3. Consider Δ_{BC}, with the A-Box statement:

Tamoxifen5YrCourse(TamE5Yr)

5 years worth of tamoxifen is a type of tamoxifen treatment, and tamoxifen treatments are a type of hormone receptor antagonist treatment. Therefore the T-Box of Δ_{BC}, contains the class of 5 years of Tamoxifen (Tamoxifen5YrCourse), which is a subclass of Tamoxifen (Tamoxifen), which is itself a subclass of hormone receptor antagonists (HormoneRARegimeType), both the following hold, even though they are not in the A-box of Δ_{BC}.

$\Delta_{BC} \vdash_{Ont}$ Tamoxifen(TamE5Yr) and

$\Delta_{BC} \vdash_{Ont}$ HormoneRARegimeType(TamE5Yr)

While we could use the logical closure of the A-Box as the basis for a set of facts in $\mathcal{G}_{\mathcal{K}}$, this creates problems if we either infer another literal as the claim of an argument or assume another literal as an additional fact to be used with the ontology. A better solution is to check for entailment as and when we need it.

Example 3.3.4. Let Δ_{BC} be the breast cancer ontology, and Women(MsJones) be a literal in $\mathcal{F}_{\mathcal{K}}$. Note, however, that Women(MsJones) is not in Δ_{BC}, but that in the breast cancer ontology, Women is a sub-class of People. Therefore $\Delta_{BC} \cup \{$Women(MsJones)$\} \vdash_{Ont}$People(MsJones).

Thus we use ontological entailment both on literals already in the the ontology and in relation to new literals.

Definition 3.3.5. Let $\phi, \psi \in \mathcal{F}_{\mathcal{K}}$ be two formulae and Δ be some DL ontology. Then the meta-predicate $conflict_{\Delta}(\phi, \psi)$, holds if all of the following three conditions hold:

1. $\Delta \cup \{\phi \cup \psi\} \vdash_{Ont} \perp$
2. $\Delta \cup \{\phi\} \nvdash_{Ont} \perp$
3. $\Delta \cup \{\psi\} \nvdash_{Ont} \perp$

Example 3.3.6. Using Δ_{BC}, we can see that Women(MsJones) and Men(MsJones) are in conflict, as Men and Women are defined as being disjoint to one another in Δ_{BC} and so $conflict_{\Delta}$(Women(MsJones), Men(MsJones)) holds.

Note that by definition 3.2.1, Δ is consistent and satisfiable (and thus $\Delta \nvdash_{Ont} \perp$). The definition of conflict is a little more complex than that of entailment, hence the use of the meta-predicate, $conflict_{\Delta}()$. However, this still uses the same idea of checking with the ontology as and when needed. Note also that conflict is symmetric i.e. $conflict_{\Delta}(\phi, \psi) \leftrightarrow conflict_{\Delta}(\psi, \phi)$.

It will also be useful to compare conflict between a formula and a set of formulae.

Definition 3.3.7. Let $\phi \in \mathcal{F}_{\mathcal{K}}$ be a formula and $\Psi \subset \mathcal{F}_{\mathcal{K}}$ some set of formulae. Let Δ be some DL ontology. Then the meta-predicate $conflict_{\Delta}(\Psi, \phi)$, holds if all of the following three conditions hold:

1. $\Delta \cup \Psi \cup \{\phi\} \vdash_{Ont} \perp$
2. $\Delta \cup \Psi \nvdash_{Ont} \perp$
3. $\Delta \cup \{\phi\} \nvdash_{Ont} \perp$

Example 3.3.8. Δ_{BC} contains the assertion TamoxifenTreatment(TamE5Yr). Therefore, if we consider the negation of this atom, and the set of all the literals in Δ_{BC}, we can see that $conflict_{\Delta}(Literals_{\mathcal{K}}(\Delta_{BC}), \neg$TamoxifenTreatment(TamE5Yr)$)$ holds.

3.4 Defeasible Reasoning in OAF

Now I have introduced the language and the use of the ontology, I can present the argumentative framework. I start by introducing the idea of a defeasible derivation, and from there move onto the construction of arguments.

Definition 3.4.1. Let $\mathcal{K}$ be a vocabulary, $\mathcal{G}_{\mathcal{K}}$ the set of literals, $\phi \in \mathcal{G}_{\mathcal{K}}$ a literal and $\Omega = (\Theta, \Delta)$ be a S-OAF. A ***defeasible derivation*** of ϕ from Ω, denoted $\Omega \vdash_{Def} \phi$, consists of a finite sequence $\phi_1, \ldots, \phi_n$ of literals, and each literal ϕ_i is in the sequence because:

1. $\Delta \cup \{\phi_1 \ldots \phi_{i-1}\} \vdash_{Ont} \phi_i$ or
2. there exists a rule, $r \in \Theta$, s.t. $Head(r) = \phi_i$ and $Body(r) = \lambda_1 \ldots \lambda_n$, and for each λ_k in $Body(r)(1 < k < n)$, λ_k is an element ϕ_j of the sequence appearing before ϕ_i (i.e $j < i$)

This definition of defeasible derivation means that there is a derivation for a literal ϕ if there is a set of literals (either in the ontology, given as extra facts or as the result of defeasible derivation) that entail ϕ, of if ϕ is the head of a rule and we have a derivation for each element of the body of the rule. Defeasible Derivation is monotonic (c.f. DeLP), but unlike DeLP, even an Ω with no facts can still have defeasible derivations (Garcia & Simari amend this in secn. 6 of their paper) as I allow rules with empty bodies. In the example below, I will assume that there is some Journal Of Medical Information (JMI) which published a study in 1978 on the effect of breast cancer on survival, and I note that the breast cancer ontology, Δ_{BC}, already contains a set of changes in risk of differing magnitude for different outcomes, including overall survival (OS).

Example 3.4.2. Let $\mathcal{K}$ be a vocabulary and $\mathcal{G}_{\mathcal{K}}$ the set of literals. Let $\Gamma \subset \mathcal{G}_{\mathcal{K}}$ be a set of literals and recall that Δ_{BC} is the breast cancer ontology. $\Omega = (\Theta, \Delta_{BC} \cup \Gamma)$ is a S-OAF where Γ and Θ are as below:

Γ ={ People(MsJones), hasDisease(MsJones,ProtoBreastCancer) }

Θ = {JMI1978}

JMI1978: People(x)
$\wedge$ hasDisease(x,y) $\wedge$ DiseaseTypes(y)
$\wedge$ DecreasedOS(z) $\wedge$ hasValue(z, 0.7)
$\Rightarrow$ hasDeltaRisk(x,z)

Then {BreastCancerTypes(ProtoBreastCancer)} $\cup$ Δ_{BC} $\vdash_{Ont}$ DiseaseTypes(ProtoBreastCancer) holds as does $\Delta_{BC} \vdash_{Ont}$ {DecreasedOS(DecreasedOS0.7), hasValue(DecreasedOS0.7, 0.7)} , and so we can satisfy the body of JMI1978, and hence $\Omega \vdash_{Def}$ hasDeltaRisk(MsJones, DecreasedOS0.7).

Note that the definition of defeasible inference ensures that every literal that is entailed by the ontology is also considered to be defeasibly entailed. However, as we shall see later, such literals do not play a part in our argumentation.

Example 3.4.3. Recall that Δ_{BC} is the breast cancer ontology. Then let Ω be a S-OAF such that Ω = $(\emptyset, \Delta_{BC})$. Then Tamoxifen5YCourse(TamE5Yr) is a literal in Δ_{BC}, and from the definition above, Ω $\vdash_{Def}$ Tamoxifen5YCourse(TamE5Yr) .

The most significant point is that we can combine defeasible and ontological reasoning.

Example 3.4.4. Let $\mathcal{K}$ be a vocabulary and $\mathcal{G}_{\mathcal{K}}$ the set of literals. Let $\Gamma \subset \mathcal{G}_{\mathcal{K}}$ be a set of literals and recall that Δ_{BC} is the breast cancer ontology. $\Omega = (\Theta, \Delta_{BC} \cup \Gamma)$ is a S-OAF where Γ and Θ are as below:

Γ ={People(MsJones), hasDisease(MsJones,ProtoBreastCancer)}

Θ = {JMP1976}

JMP1976: People(x)

$\wedge$ hasDisease(x,y) $\wedge$ BreastCancerTypes(y)

$\wedge$ Tamoxifen5YCourse(z)

$\Rightarrow$ hasTreatment(x,z)

In this case, $\Omega \vdash_{Def}$ hasTreatment(MsJones, TamE5Yr). However, it is also the case that the ontology defines a class for those who take tamoxifen (PeopleOnTamoxifen). Therefore it is also the case that $\Omega \vdash_{Def}$ PeopleOnTamoxifen(MsJones).

Definition 3.4.5. Let Δ be a DL ontology, ϕ be a formula in $\mathcal{F}_{\mathcal{K}}$ and Θ be a set of defeasible rules. Then Ω is a S-OAF such that Ω = (Θ, Δ). There is a ***strict derivation*** of ϕ, denoted $\Omega \vdash_{Ont} \phi$, iff $\Delta \vdash_{Ont} \phi$ holds.

Example 3.4.6. In 3.4.4 above, the derivation of Tamoxifen5YCourse(TamE5Yr) is strict; that of hasTreatment(MsJones, TamE5Yr) is defeasible. The derivation of PeopleOnTamoxifen(MsJones) is also defeasible.

Definition 3.4.7. Let Θ be a set of defeasible rules and Δ a DL ontology. Then Ω is an S-OAF such that Ω = (Θ, Δ). A set of rules Θ is ***contradictory*** with respect to Δ iff there exist two formulae, $\phi_1$and ϕ_2 such that $\Omega \vdash_{Def} \phi_1$ and $\Omega \vdash_{Def} \phi_2$ and $conflict_{\Delta}(\phi_1, \phi_2)$ holds.

Example 3.4.8. Let $\mathcal{K}$ be a vocabulary and $\mathcal{G}_{\mathcal{K}}$ the set of literals. Let $\Gamma \subset \mathcal{G}_{\mathcal{K}}$ be a set of literals and recall that Δ_{BC} is the breast cancer ontology. $\Omega = (\Theta, \Delta_{BC} \cup \Gamma)$ is a S-OAF where Γ and Θ are as below:

Γ ={Woman(MsJones), hasDisease(MsJones,ProtoBreastCancer)}

Θ = {JMI1978, JWS1988}

and

JMI1978: People(x)

$\wedge$ hasDisease(x,y) $\wedge$ DiseaseTypes(y)

$\wedge$ DecreasedOS(z) $\wedge$ hasValue(z, 0.7)

$\Rightarrow$ hasDeltaRisk(x,z)

JWS1988: Women(x)

$\wedge$ IncreasedOS(y) $\wedge$ hasValue(y, 1.1)

$\Rightarrow$ hasDeltaRisk(x,y)

Then we can see that the following hold:

$\Omega \vdash_{Def}$ hasDeltaRisk(MsJones, IncreasedOS1.1)

$\Omega \vdash_{Def}$ hasDeltaRisk(MsJones, DecreasedOS0.7)

$conflict_{\Delta_{BC}}$((hasDeltaRisk(MsJones, DecreasedOS0.7), (hasDeltaRisk(MsJones, IncreasedOS1.1))

and hence Θ is contradictory with respect to Δ_{BC}.

Note that in the example above, since the rules refer to overall survival, I do not include any reference to a specific disease.

3.5 Defeasible Argumentation

I have introduced the definitions of defeasible and strict derivation above. I now use these derivations to define a defeasible argument. An argument is a pair, $\langle A, \phi \rangle$ where A provides a minimal, consistent derivation for ϕ given the facts in the ontology, Δ.

Definition 3.5.1. Let $\mathcal{K}$ be a DL ontology language, $\mathcal{O}_{\mathcal{K}}$ the set of formulae in $\mathcal{K}$, and $\Delta \subseteq \mathcal{O}_{\mathcal{K}}$ an ontology. Let $\mathcal{F}_{\mathcal{K}}$ be the set of formulae in $\mathcal{L}_{\mathcal{K}}$ and $\mathcal{R}_{\mathcal{K}}$ be the associated set of defeasible rules. Let $\Theta \subseteq \mathcal{R}_{\mathcal{K}}$ be a set of defeasible rules and $\Gamma \subseteq \mathcal{F}_{\mathcal{K}}$ a set of facts. Let $\phi \in \mathcal{G}_{\mathcal{K}}$ be a literal, and Ω_{Δ} be a S-OAF, such that $\Omega_{\Delta} = (\Theta, \Delta)$. $\langle A, \phi \rangle$ is an ***argument*** for ϕ with respect to Δ, if $A \subseteq \Theta$, and the following conditions hold:

1. $(A, \Delta) \vdash_{Def} \phi$

2. A is not contradictory with respect to Δ

3. A is minimal: there is no proper subset *A'* of *A* such that *A'* satisfies condition (1).

The set of all arguments with respect to the vocabulary $\mathcal{K}$ is denoted $\mathcal{A}_{\mathcal{K}}$ and with respect to an ontology Δ is denoted $\mathcal{A}_{\Delta}$.

In short, an argument $\langle A, \phi \rangle$ is a minimal non-contradictory set of defeasible rules, or their labels, that provides a defeasible derivation for a given literal, ϕ . A will be called the ***support*** and ϕ will also be called the ***claim*** supported by *A*.

Example 3.5.2. Let $\mathcal{K}$ be a vocabulary and $\mathcal{G}_{\mathcal{K}}$ the set of literals. Let $\Gamma \subset \mathcal{G}_{\mathcal{K}}$ be a set of literals and recall that Δ_{BC} is the breast cancer ontology. $\Omega = (\Theta, \Delta_{BC} \cup \Gamma)$ is a S-OAF where Γ and Θ are as below:

Γ = {Woman(MsJones), hasDisease(MsJones,ProtoBreastCancer)}

Θ = {JPS1978, JWS1988}

where

JPS1978: People(x)

$\wedge$ hasDisease(x,y) $\wedge$ DiseaseTypes(y)

$\wedge$ DecreasedOS(z) $\wedge$ hasValue(z, 0.7)

$\Rightarrow$ hasDeltaRisk(x,z)

JWS1988: Women(x)

$\wedge$ IncreasedOS(y) $\wedge$ hasValue(y, 1.1)

$\Rightarrow$ hasDeltaRisk(x,y)

and we recall that

$\Delta_{BC} \cup$ {Woman(MsJones)} $\vdash_{Ont}$ People(MsJones))

$\Delta_{BC} \vdash_{Ont}$ DiseaseTypes(ProtoBreastCancer))

and that there are explicit changes in risk in the A-box of Δ_{BC}:

DecreasedOS(DecreasedOS0.7) $\wedge$ hasValue (DecreasedOS0.7, 0.7)

IncreasedOS(IncreasedOS1.1) $\wedge$ hasValue (IncreasedOS1.1, 1.1)

then we can form the following arguments:

a_1 $\langle${JPS1978}, hasDeltaRisk(MsJones, DecreasedOS0.7)$\rangle$

a_2 $\langle${JWS1988}, hasDeltaRisk(MsJones, IncreasedOS1.1)$\rangle$

and $conflict_{\Delta}$(hasDeltaRisk(MsJones, DecreasedOS0.7), hasDeltaRisk(MsJones, IncreasedOS1.1)) holds.

As in DeLP, the construction of arguments is non-monotonic, as illustrated in the next example. More generally, given two ontologies where $\Delta \subset \Delta'$, it is not the case that $\mathcal{A}_{\Delta} \subset \mathcal{A}_{\Delta'}$. This is because adding assertions to an ontology can invalidate arguments that would have been considered acceptable before the assertions were added.

Example 3.5.3. Consider the example above. There are two arguments:

a_1 $\langle${JPS1978}, hasDeltaRisk(MsJones, DecreasedOS0.7)$\rangle$

a_2 $\langle${JWS1988}, hasDeltaRisk(MsJones, IncreasedOS1.1)$\rangle$

However, if we add the literal hasDeltaRisk(MsJones, IncreasedOS1.1) to Δ_{BC}, then (Δ_{BC} $\cup$ hasDeltaRisk(MsJones, IncreasedOS1.1)) $\cup$ hasDeltaRisk(MsJones, DecreasedOS0.7) $\vdash_{Ont} \perp$. Recalling that one of the conditions of an argument is that its support is a non-conflicting set of rules, we can see that a_2 is no longer an argument. More generally, this example demonstrates the primacy that DeLP and OAF place on facts: it is simply not possible to construct an argument whose claim conflicts with a fact. One side-effect of this is that if we alter our facts to be defeasibly asserted unconditionals (i.e. instead of $\Delta \vdash_{Ont} \phi$ for some literal ϕ, we instead have a rule of the form $r :\Rightarrow \phi$), then this problem disappears, and argument construction is monotonic.

It will sometimes be useful to refer to the elements of an argument.

Definition 3.5.4. The function $Support$ returns the support of an argument $\langle A,\, \phi\rangle$, while the function $Claim$ returns the claim i.e. $Support(\langle A, \phi\rangle)$ = A and $Claim(\langle A, \phi\rangle)$ = ϕ.

It is also helpful to be able to describe the size of an argument:

Definition 3.5.5. Let $\langle A,\, \phi\rangle$ be an argument in $\mathcal{A}_\Delta$. Then the ***size*** of an argument is defined as the number of elements in A, and is returned by the function $Size(\langle A,\, \phi\rangle)$: $\mathcal{A}_\Delta \mapsto I\!N$.

I also define the idea of a subargument

Definition 3.5.6. An argument $\langle B, \psi\rangle$ is a ***sub-argument*** of $\langle A, \phi\rangle$ if $B \subseteq A$ and a ***proper sub-argument*** if $B \subset A$.

3.5.1 Counter-arguments and Attacks

Definition 3.5.7. Let $\mathcal{K}$ be a vocabulary, and $\phi_1 ... \phi_n$ literals in $\mathcal{G}_\mathcal{K}$. $\langle A_1, \phi_1\rangle$ ***attacks*** $\langle A_2, \phi_2\rangle$ at literal ϕ iff there exists a sub-argument $\langle A, \phi\rangle$ of $\langle A_2, \phi_2\rangle$ such that $conflict_\Delta(\phi, \phi_1)$ holds. Attacks are described as direct iff $\langle A, \phi\rangle = \langle A_2, \phi_2\rangle$, and indirect if $\langle A, \phi\rangle$ is a proper sub-argument of $\langle A_2, \phi_2\rangle$. One argument that attacks another is known as a ***counter-argument*** to it. The claim of an argument to which there is a counter-argument is described as being ***rebutted***.

Example 3.5.8. I reuse example 3.5.2. Let $\mathcal{K}$ be a vocabulary and $\mathcal{G}_\mathcal{K}$ the set of literals. Let $\Gamma \subset \mathcal{G}_\mathcal{K}$ be a set of literals and recall that Δ_{BC} is the breast cancer ontology. $\Omega = (\Theta, \Delta_{BC} \cup \Gamma)$ is a S-OAF where Γ and Θ are as below:

Γ = {Woman(MsJones), hasDisease(MsJones,ProtoBreastCancer)}

Θ = {JPS1978, JWS1988}

where

JPS1978: People(x)

$\wedge$ hasDisease(x,y) $\wedge$ DiseaseTypes(y)

$\wedge$ DecreasedOS(z) $\wedge$ hasValue(z, 0.7)

$\Rightarrow$ hasDeltaRisk(x,z)

JWS1988: Women(x)

$\wedge$ IncreasedOS(y) $\wedge$ hasValue(y, 1.1)

$\Rightarrow$ hasDeltaRisk(x,y)

and we recall that

$\Delta_{BC} \cup$ Woman(MsJones) $\vdash_{Ont}$ People(MsJones))

$\Delta_{BC} \vdash_{Ont}$ DiseaseTypes(ProtoBreastCancer))

then we can form the following arguments:

a_1 $\langle${JPS1978}, hasDeltaRisk(MsJones, DecreasedOS0.7)$\rangle$

a_2 $\langle${JWS1988}, hasDeltaRisk(MsJones, IncreasedOS1.1)$\rangle$

and we recall that $conflict_\Delta$(hasDeltaRisk(MsJones, DecreasedOS0.7), hasDeltaRisk(MsJones, IncreasedOS1.1)) holds. Therefore a_1 attacks a_2.

Recall that we defeasibly entail any literal that is strictly entailed. It may appear that this would cause problems with our argumentation, as we could have a mixture of strictly and defeasibly entailed literals as the claims of arguments; however, this is not the case, as the following two proofs show.

Proposition 3.5.9. *No counter-argument exists for an argument* $\langle \emptyset, \phi \rangle$.

Proof. If there exists an argument $\langle \emptyset, \phi_1 \rangle$, then from the definition of an argument, it is the case that $(\emptyset, \Delta) \vdash_{Def} \phi_1$. Therefore there is a strict entailment for ϕ_1 from Δ (i.e. $\Delta \vdash_{Ont} \phi_1$). Suppose that there exists a counter-argument $\langle A, \phi_2 \rangle$ for $\langle \emptyset, \phi_1 \rangle$. Then from 3.5.7 and 3.4.7, $(A \cup \emptyset)$ should be a contradictory set of defeasible rules with respect to Δ, and hence A should also be contradictory w.r.t Δ. However, this contradicts our definition of an argument, which requires that the support of an argument, A should not be contradictory. Thus, a strictly entailed literal cannot be rebutted by one that is only defeasible entailed . $\square$

Proposition 3.5.10. *The argument* $\langle \emptyset, \psi \rangle$ *cannot be a counter-argument for any argument* $\langle A, \phi \rangle$.

Proof. A similar technique is used as for the proof above: Suppose that $\langle \emptyset, \phi_1 \rangle$ is a counter-argument for $\langle A, \phi_2 \rangle$, then $(A \cup \emptyset)$ should be a contradictory set of rules with respect to Δ. Therefore A should be a contradictory set with respect to Δ. However, this contradicts our definition of an argument. Thus, arguments for defeasibly entailed literals cannot be attacked by literals that are strictly entailed. $\square$

These two results ensure that strict entailment only takes part in defeasible derivation of literals, rather than in the rebuttal of arguments. I can also show that no argument attacks itself.

Proposition 3.5.11. *No argument can attack itself*

Proof. Assume $\langle A, \phi \rangle$ attacks itself. Then there should exist a counter-argument point ψ in A such that $conflict_\Delta(\psi, \phi)$ is true. Therefore A would be contradictory wrt Δ, and so $\langle A, \phi \rangle$ would not be an argument. $\square$

3.5.2 Comparing Arguments using Rule priorities

To resolve conflict between arguments, I use information about the rules in the support of the argument, based on a preference ordering between the rules.

Definition 3.5.12. Let ">" be a ***preference relationship*** defined among defeasible rules. Given two arguments $\langle A_1, \phi_1 \rangle$ and $\langle A_2, \phi_2 \rangle$, the argument $\langle A_1, \phi_1 \rangle$ is preferred over $\langle A_2, \phi_2 \rangle$, denoted $\langle A_1, \phi_1 \rangle > \langle A_2, \phi_2 \rangle$ iff:

1. There exists a rule $r_a \in A_1$, and a rule $r_b \in A_2$, such that $r_a > r_b$ and
2. There is no $r'_b \in A_2$ and $r'_a \in A_1$ such that $r'_b > r'_a$

Example 3.5.13. Let us consider the S-OAF defined in 3.5.2, which resulted in the following arguments:

$a_1\langle$ {JPS1978}, hasDeltaRisk(MsJones, DecreasedOS0.7) $\rangle$

$a_2\langle$ {JWS1988}, hasDeltaRisk(MsJones, IncreasedOS1.1) $\rangle$

I prefer rules based on the nature of the date of their source, (preferring more recent rules - a formal treatment of this is left until chapter 6). The dates of the sources are clear from the labels and we therefore have a preference ordering of: JWS1988 > JPS1978. Since a_1 and a_2 attack each other, we conclude that a_2 is preferred to a_1.

Note that because no argument attacks itself, no argument can defeat itself, and so OAF does not have self-defeating arguments.

3.6 Defeaters and Argumentation Lines

I have now shown how we can use a combination of ontological and defeasible inference to develop arguments. One of our expectations is that these arguments will conflict, and I now present a technique for calculating the defeat status of arguments that are attacked.

Definition 3.6.1. Let $\langle A_1, \phi_1 \rangle$ and $\langle A_2, \phi_2 \rangle$ be two arguments. $\langle A_1, \phi_1 \rangle$ is a ***proper defeater*** for $\langle A_2, \phi_2 \rangle$ at literal ϕ iff there exists a sub-argument $\langle A, \phi \rangle$ of $\langle A_2, \phi_2 \rangle$ such that $\langle A_1, \phi_1 \rangle$ attacks $\langle A_2, \phi_2 \rangle$ at ϕ and $\langle A_1, \phi_1 \rangle$ > $\langle A, \phi \rangle$.

Of course, it may be the case that the two arguments are unrelated by a preference ordering

Definition 3.6.2. Let $\langle A_1, \phi_1 \rangle$ and $\langle A_2, \phi_2 \rangle$ be two arguments. $\langle A_1, \phi_1 \rangle$ is a ***blocking defeater*** for $\langle A_2, \phi_2 \rangle$ at ϕ iff there exists a sub-argument $\langle A, \phi \rangle$ of $\langle A, \phi_2 \rangle$ such that $\langle A_1, \phi_1 \rangle$ attacks $\langle A_2, \phi_2 \rangle$ at ϕ and $\langle A_1, \phi_1 \rangle$ is unrelated to $\langle A_2, \phi_2 \rangle$ by a preference-ordering to $\langle A, \phi \rangle$ i.e. $\langle A_1, \phi_1 \rangle \not> \langle A, \phi \rangle$ and $\langle A, \phi \rangle \not> \langle A_1, \phi_1 \rangle$.

The argument $\langle A_1, \phi_1 \rangle$ is a ***defeater*** for $\langle A_2, \phi_2 \rangle$ iff it is either a proper or a blocking defeater.

Example 3.6.3. Using examples 3.5.2 and 3.5.13, we recall that there are two arguments:

$a_1\langle$ {JPS1978}, hasDeltaRisk(MsJones, DecreasedOS0.7) $\rangle$

$a_2\langle$ {JWS1988}, hasDeltaRisk(MsJones, IncreasedOS1.1) $\rangle$

with the following preference order between rules:

JWS1988 > JPS1978

we can see that since a_2 attacks a_1 and JWS1988 is preferred to JPS1978, a_2 is a proper defeater of a_1.

Definition 3.6.4. Let $\langle A_1, \phi_1\rangle$ be a defeater for $\langle A_2, \phi_2\rangle$, such that $\langle A_1, \phi_1\rangle$ counter-argues $\langle A_2, \phi_2\rangle$ at a sub-argument, $\langle A, \phi\rangle$. $\langle A, \phi\rangle$ is termed the ***disagreement sub-argument***

This use of the word disagreement is consistent with the terminology used by Garcia & Simari [32] . Other authors use the term disputed arguments to capture the same intuition, but I shall continue with this usage here.

Example 3.6.5. Let $\mathcal{K}$ be a vocabulary and $\mathcal{G}_{\mathcal{K}}$ the set of literals. Let $\Gamma \subseteq \mathcal{G}_{\mathcal{K}}$ be a set of literals and recall that Δ_{BC} is the breast cancer ontology. $\Omega = (\Theta, \Delta_{BC} \cup \Gamma)$ is a S-OAF where Γ and Θ are as below:

Γ = {Woman(MsJones), hasDisease(MsJones,ProtoAdenoCaBreast)}

Θ = {NEJM2003, BJC1997, BMJ2001, Lancet2001}

NEJM2003: People(x)
$\wedge$ hasDisease(x,y) $\wedge$ BreastCancerTypes(y)
$\wedge$ Tamoxifen5YrCourse(z)
$\Rightarrow$ hasTreatment(x,z)

BJC1997: People(x)
$\wedge$ hasDisease(x,y) $\wedge$ BreastCancerTypes(y)
$\wedge$ Tamoxifen2YrCourse(z)
$\Rightarrow$ hasTreatment(x,z)

BMJ2001: People(x)
$\wedge$ hasDisease(x,y) $\wedge$ BreastCancerTypes(y)
$\wedge$ ECMFRegimeType(z)
$\Rightarrow$ hasTreatment(x,z)

Lancet2001: People(x)
$\wedge$ hasTreatment(x,y) $\wedge$ Tamoxifen(y)
$\wedge$ IncreasedRiskEndometrialCancer(z) $\wedge$ hasValue(z, 2.3)
$\Rightarrow$ hasDeltaRisk(x,z)

We form the following arguments from Ω:

$a_1\langle$ {NEJM2003}, hasTreatment(MsJones,TamE5yr) $\rangle$

$a_2\langle$ {BJC1997}, hasTreatment(MsJones, TamD2Yr) $\rangle$

$a_3\langle$ {BMJ2001}, hasTreatment(MsJones, ECMFTypeInstance) $\rangle$

$a_4\langle$ {NEJM2003, Lancet2001}, hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer2.3)$\rangle$

$a_5\langle$ {BJC1997, Lancet2001}, hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer2.3)$\rangle$

We also note that

$conflict_{\Delta_{BC}}$(hasTreatment(MsJones, TamE5yr), hasTreatment(MsJones, TamD2yr))

$conflict_{\Delta_{BC}}$(hasTreatment(MsJones, TamE5yr), hasTreatment(MsJones, ECMFTypeInstance))

$conflict_{\Delta_{BC}}$(hasTreatment(MsJones, TamD2yr), hasTreatment(MsJones,ECMFTypeInstance))

hold

In the example above, if NEJM2003 > BMJ2001 > BJC1997 > Lancet2001, then a_1 is a proper defeater for a_2, a_3 and a_5. Furthermore, in the attack of a_1 on a_5, a_3 is the disagreement sub-argument, and demonstrates the value of the disagreement sub-argument, in that it lets us concentrate on *where* two arguments disagree, rather than simply saying that they do.

Definition 3.6.6. Let Θ be a set of defeasible rules and Δ be a DL ontology. Let Ω be a S-OAF such that Ω= (Θ, Δ). An ***argumentation line*** for $\langle A_0, \phi_0\rangle$ is a sequence of arguments from Ω, denoted $\Lambda = [\langle A_0, \phi_0\rangle, \langle A_1, \phi_1\rangle ....]$ where each element in the sequence $\langle A_i, \phi_i\rangle, i > 0$, is a defeater of its predecessor, $\langle A_{i-1}, \phi_{i-1}\rangle$.

Example 3.6.7. Using the Ω as described in 3.6.5, and allowing NEJM2003 > BMJ2001 > BJC1997 > Lancet2001, then considering argument a_5, one argument line would be:

Λ = [a_5, a_3]

An argumentation line is therefore a sequence of arguments, which alternate in their support and defeat of the initial element. It is useful to be able to distinguish these two different sets.

Definition 3.6.8. Let $\Lambda = [\langle A_0, \phi_0\rangle, \langle A_1, \phi_1\rangle, \ldots]$ be an argumentation line. I define the set of ***supporting arguments*** $\Lambda_S = \{\langle A_0, \phi_0\rangle, \langle A_2, \phi_2\rangle, \ldots\}$ and the set of ***interfering arguments***, $\Lambda_I = \{\langle A_1, \phi_1\rangle, \langle A_3, \phi_3\rangle, \ldots\}$.

For any given argument, $\langle A_0, \phi_0\rangle$, there can be many defeaters, each of which generates its own argument line.

Example 3.6.9. Using the same S-OAF as in 3.6.5, and again allowing NEJM2003 > BMJ2001 > BJC1997 > Lancet2001, we see that one argumentation line is:

Λ = [a_5, a_3]

which we can divide into two sets, Λ_S and Λ_I:

Λ_S = $\{a_5\}$

Λ_I = $\{a_3\}$

Definition 3.6.10. Two arguments $\langle A_1, \phi_1 \rangle$ and $\langle A_2, \phi_2 \rangle$ are ***concordant*** with respect to Δ if $A_1 \cup A_2$ is not contradictory with respect to Δ.

The notion of concordant arguments lets us enforce a restriction that each argumentation line must be consistent; we cannot have an argument that defends the root argument while at the same time attacking it.

Definition 3.6.11. Let $\Lambda = [\langle A_0, \phi_0 \rangle, \langle A_1, \phi_1 \rangle ....]$ be an argumentation line. and Δ a DL ontology. Λ is an ***acceptable argumentation line*** iff:

1. Λ is a finite sequence
2. Every pair of arguments in the set Λ_S of supporting arguments is concordant with respect to Δ and every pair of arguments in the set Λ_I of interfering arguments is concordant with respect to Δ .
3. No argument $\langle A_k, \phi_k \rangle$ in Λ is a sub-argument of an argument $\langle A_i, \phi_i \rangle$ appearing earlier in Λ $(i < k)$.
4. For all i, such that the argument $\langle A_i, \phi_i \rangle$ is a blocking defeater for $\langle A_{i-1}, \phi_{i-1} \rangle$, if $\langle A_{i+1}, \phi_{i+1} \rangle$ exists, then $\langle A_{i+1}, \phi_{i+1} \rangle$ is a proper defeater for $\langle A_{i,} \phi_i \rangle$.

Example 3.6.12. Consider the same example as in 3.6.9:

$\Lambda_S = \{a_5\}$

$\Lambda_I = \{a_3\}$

Note that the argument line Λ_1 does not include all the arguments in Ω. Specifically, it excludes a_1:

a_1 $\langle${NEJM2003}, hasTreatment(MsJones,TamE5yr)$\rangle$

even though a_1 attacks the root argument, a_5. This is because a_1 and a_3 conflict, and thus are not concordant, and so its inclusion would make the argumentation line unacceptable.

3.7 Warrant through Dialectical Analysis

Now we have suitable definitions of arguments and argument lines, I use them to create dialectical trees.

Definition 3.7.1. Let $\langle A_0, \phi_0 \rangle$ be an argument from Ω. A ***dialectical tree*** for$\langle A_0, \phi_0 \rangle$, denoted $\mathcal{T}_{\langle A_0, \phi_0 \rangle}$is defined as follows:

1. The root of the tree is labeled with $\langle A_0, \phi_0 \rangle$

2. Let N be a non-root node of the tree labeled $\langle A_n, \phi_n \rangle$ and $\Lambda = [\langle A_0, \phi_0 \rangle, \langle A_1, \phi_1 \rangle ....]$ the sequence of labels of the path from the root to N. Let $\langle B_0, \psi_0 \rangle, \langle B_1, \psi_1 \rangle ..... \langle B_k, \psi_k \rangle$ be all the defeaters for $\langle A_n, \phi_n \rangle$. For each defeater $\langle B_i, \psi_i \rangle$ $(1 \leq i \leq k)$ such that, the argumentation line $\Lambda' = [\langle A_0, \phi_0 \rangle, \langle A, \phi_1 \rangle .... \langle A_n, \phi_n \rangle, \langle B_i, \psi_i \rangle]$ is acceptable, then the node N has a child N_i labeled $\langle B_i, \psi_i \rangle$. If there is no defeater for $\langle A_n, \psi_n \rangle$ or there is no $\langle B_i, \psi_i \rangle$ such that Λ' is acceptable, then N is a leaf.

A subtree of a dialectical tree (i.e., a node with all its descendants) is not always a dialectical tree. Suppose we build an acceptable argumentation line where a defeater $\langle A, \phi \rangle$ will not be included because it would make the line unacceptable, then there might be a subsequence of the mentioned line where the same defeater could be included.

Algorithm 3.7.2. *Let $\mathcal{T}_{\langle A,\phi \rangle}$ be a dialectical tree for $\langle A, \phi \rangle$. The corresponding marked dialectical tree, denoted $\mathcal{T}^*_{\langle A,\phi \rangle}$, will be obtained marking every node in $\mathcal{T}_{\langle A,\phi \rangle}$ as follows:*

1. *All leaves in $\mathcal{T}_{\langle A,\phi \rangle}$ are marked as "U" in $\mathcal{T}^*_{\langle A,\phi \rangle}$.*
2. *Let $\langle B, \psi \rangle$ be an inner node of $\mathcal{T}_{\langle A,\phi \rangle}$. Then $\langle B, \psi \rangle$ will be marked "U" in $\mathcal{T}^*_{\langle A,\phi \rangle}$ if every child of $\langle B, \psi \rangle$ is marked as "D". The node $\langle B, \psi \rangle$ will be marked as "D" in $\mathcal{T}^*_{\langle A,\phi \rangle}$ iff it has at least one child marked as "U".*

Definition 3.7.3. Let $\langle A, \phi \rangle$ be an argument and $\mathcal{T}^*_{\langle A,\phi \rangle}$ its associated marked dialectical tree. Then ϕ is ***warranted*** iff the root of $\mathcal{T}^*_{\langle A,\phi \rangle}$ is marked as "U". I will say that A is a ***warrant*** for ϕ. An argument $\langle B, \psi \rangle$ where $\mathcal{T}^*_{\langle B,\psi \rangle}$ is marked as "D" is described as ***unwarranted***.

This definition means that the warrant status of an argument is decided by the marked dialectical tree constructed for that argument.

Proposition 3.7.4. *If ϕ has a strict derivation from Ω, then ϕ is warranted.*

One important difference from DeLP is that whereas they work in a goal-directed manner by posing queries and then constructing arguments, I work in a forward-chaining manner, seeking to draw inferences from the our data. As a result, OAF generates a great number of arguments, for different outcomes, and then resolve them, rather than constructing arguments to satisfy queries. Therefore we need some simple techniques to filter the arguments that we consider so that we can concentrate on ones that we consider significant.

Example 3.7.5. Let $\mathcal{K}$ be a vocabulary and $\mathcal{G}_{\mathcal{K}}$ the set of literals. Let $\Gamma \subset \mathcal{G}_{\mathcal{K}}$ be a set of literals and recall that Δ_{BC} is the breast cancer ontology. $\Omega = (\Theta, \Delta_{BC} \cup \Gamma)$ is a S-OAF where Γ and Θ are as below:

Γ ={ People(MsJones), hasDisease(MsJones,ProtoBreastCancer) }

Θ = {JMI1976}

JMI1976: People(x)
$\wedge$ hasDisease(x,y) $\wedge$ BreastCancerTypes(y)
$\wedge$ Tamoxifen5YrCourse(TamE5Yr)
$\Rightarrow$ hasTreatment(x, TamE5Yr)

In this case, $\Omega \vdash_{Def}$ hasTreatment(MsJones, TamE5Yr). However, it is also the case that the ontology defines a class for those who take tamoxifen, PeopleOnTamoxifen. Therefore it is also the case that $\Omega \vdash_{Def}$ PeopleOnTamoxifen(MsJones), and so there are two arguments:

a_1 $\langle$JMI1976, hasTreatment(MsJones, TamE5Yr)$\rangle$

a_2 $\langle$JMI1976, PeopleOnTamoxifen(MsJones)$\rangle$

Note that the support of the two arguments is the same

In this case, the first argument is of more interest to us than the second, which is simply the strict derivation as a result of the first. The knowledge in the ontology is about the structure of our domain. It is important, but for people in the domain, it is regarded as being 'obvious'. The breast cancer ontology contains information about types of breast cancer, treatments, staging, etc. While this is knowledge, from a clinical aspect it is 'background' knowledge. In contrast, the knowledge captured in the defeasible rules is considered important, and debatable, even within the domain. I therefore refer to this sort of information as being 'foreground' knowledge, and I consider such information to be more interesting than the background knowledge. Therefore, when we come to filter arguments, we shall concentrate on arguments whose claims are defeasible, which I refer to as ***foreground*** arguments. This appear to represent a mixing of domain-specific and computational aspects. However, irrespective of the domain, I suggest that arguments whose claims are the result of ontological reasoning (background arguments) are, in that domain, of less interest than foreground arguments. Therefore, although the specific definition of an argument as foreground or background depends on the domain-specific ontology, the principle of the distinction is domain independent.

Definition 3.7.6. Let $\mathcal{K}$ be a vocabulary, Δ be a DL ontology and $\mathcal{A}_\Delta$ the set of arguments. Let $r \in \mathcal{R}_\mathcal{K}$ be a defeasible rule. Then an argument $\langle A_1, \phi_1\rangle \in \mathcal{A}_\Delta$ is a ***foreground*** argument iff $\langle A_1, \phi_1\rangle$ is an argument and ϕ_1 is the head of a defeasible rule in A_1.

Example 3.7.7. I reuse example 3.7.5 from above. Let $\mathcal{K}$ be a vocabulary and $\mathcal{G}_\mathcal{K}$ the set of literals. Let $\Gamma \subset \mathcal{G}_\mathcal{K}$ be a set of literals and recall that Δ_{BC} is the breast cancer ontology. $\Omega = (\Theta, \Delta_{BC} \cup \Gamma)$ is a S-OAF where Γ and Θ are as below:

Γ ={Women(MsJones), hasDisease(MsJones,ProtoBreastCancer)}

Θ = {JMI1976}

JMI1976: People(x)
$\wedge$ hasDisease(x,y) $\wedge$ BreastCancerTypes(y)
$\wedge$ Tamoxifen5YrCourse(TamE5Yr)
$\Rightarrow$ hasTreatment(x, TamE5Yr)

Then we can form an argument:

a_1 ⟨{JMI1976}, hasTreatment(MsJones, TamE5Yr)⟩

Since hasTreatment(MsJones, TamE5Yr) = $Head$(JMI1976), a_1 is a foreground argument. There is another argument:

a_2 ⟨ {JMI1976}, WomenOnTamoxifen(MsJones) ⟩

Where a_2 is a background argument.

3.8 Using OAF

The intention behind OAF is to help us reason about medical conditions and treatments using argumentation. As such we need to construct our arguments and then resolve them (in dialectical trees), rather than posing queries and then trying to generate answers as happens in DeLP. Obviously this is far more computationally expensive, but is necessary to satisfy our need to highlight unexpected consequences of treatment (as suggested in point (2) of the desiderata given in chapter 1:) arguments about beliefs are not only beliefs about what might happen without our intervention but also what might happen as a result of our intervention) For example, if we were aware of some side-effect in order to query whether it might happen, in many cases we would not need the system. Because of this, we start by taking all the arguments in $\mathcal{A}_\Delta$ and attempting to construct a dialectical tree for each in turn. This then allows us to see which arguments, and hence which claims, are justified.

Of course, claims that are held to be either warranted or unwarranted can still be used in constructing further arguments, although if unjustified claims are used in the defeasible derivation of a new argument, it will be attacked and found to be unwarranted in the same way as the the initial (unwarranted) argument. However, at some point we will need to decide on a set of warranted arguments, and hence claims, probably by including human intervention (see chapter 7 for more details). At this point, we need to allow these warranted inferences to be accessible to the more general database represented by the DL ontology.

The way to do this is to map our warranted claims back into the DL ontology. For practical reasons we might want to do something like adding some annotation that recorded their source as being the result of argumentation but the principle is clear - without having at least the ability to map our claims back into the DL ontology, they remain inaccessible for other uses of the database. However, because of the range of practical and technical problems that this asks, I shall not cover this in this thesis - we realise that it will be necessary in some form, without exploring the details of the procedure. The fact that the claims of our arguments are all ground literals is consistent with our intuitions about OAF arguing about individuals and their place in the world (i.e. the A-Box) not the structure of the world itself (i.e. T-Box reasoning).

3.9 Where to from here ?

This concludes my initial presentation of the OAF framework with initial modifications; I now move on to consider some refinements. Although the logic has not gone far beyond DeLP, the spirit is quite different, with a strong emphasis on using the ontology to provide a framework and set of facts with which we can argue. The work in the next chapters is divided as follows: the incorporation of valuation of outcomes into the formation of rules (Chapter 4); the development of hypothetical arguments about possible courses of treatment (Chapter 5); a prototype implementation and some simple complexity results (Chapter 6); a clinical case study (Chapter 7).

The diagrams show how I use the work in this chapter, which is itself an extension of DeLP, as the core of the argumentation formalism. In later chapters I address particular weaknesses of the current approach to develop new approaches to argument construction. Throughout the thesis, however, I use the techniques described in this chapter to evaluate and construct these new arguments in the same way as existing ones. This chapter has introduced the use of an ontology and defeasible rules to construct arguments whose claims are resolved by dialectical tree construction.

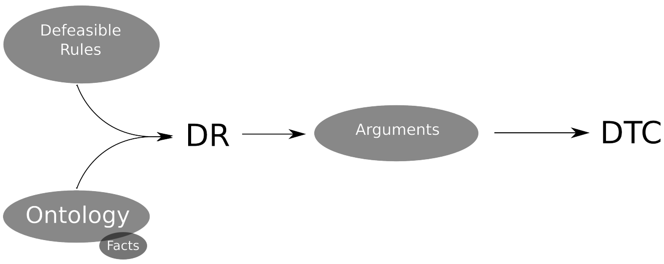

Figure 3.1: Reasoning in a Simple OAF (Chapter 3). ***DR:*** Defeasible Reasoning ***DTC:*** Dialectical Tree Construction. This is then used as the basis for reasoning in later chapters.

The next chapter introduces the use of a patient's values to rewrite defeasible rules, which are then used to construct arguments in the same way as in this chapter.

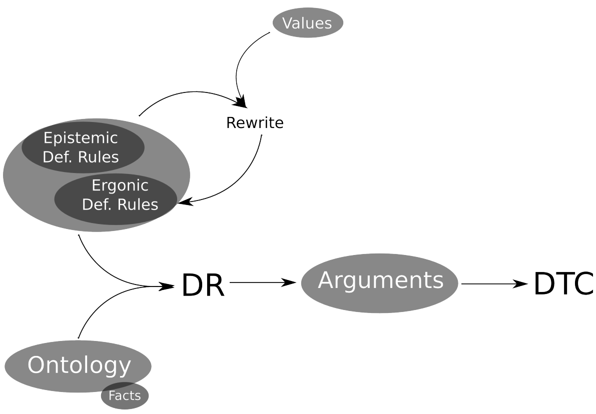

Figure 3.2: Reasoning in a Valued OAF (Chapter 4). ***DR:*** Defeasible Reasoning ***DTC:*** Dialectical Tree Construction. Note how the core structure of the Simple OAF remains, with the addition of a technique to rewrite rules using a patient's values.

Chapter 5 introduces a technique for extending arguments using a particular class of defeasible rule.

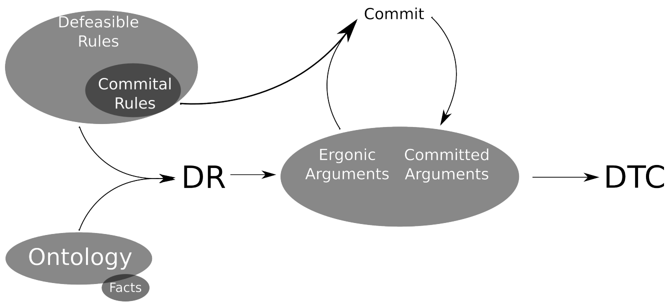

Figure 3.3: Reasoning in a Hypothetical OAF (Chapter 5). ***DR:*** Defeasible Reasoning ***DTC:*** Dialectical Tree Construction. Note how this uses the basic structure of a Simple OAF, and extends it by defining a new type of defeasible rule, Commital rules. These are used to take arguments constructed by the standard defeasible reasoning technique and develop a new type of argument. Note that the basic steps of defeasible reasoning and dialectical tree construction remain unchanged.

Chapter 4

Value-Based Argumentation

Summary

This chapter presents a technique for generating arguments for action, based on the expected outcomes of the action. I start with a short discussion of intention and patient values and how we can represent these in OAF. I show that the approach taken here bears a close relationship to existing work, and I present the central piece of work of this chapter, a function for rewriting epistemic rules as ergonic ones. These are used to generate arguments about possible actions, but at the end of the chapter I show that this only meets some of the requirements outlined in chapter 1, and the work in this chapter is the foundation for the development of hypothetical reasoning in the next. Some theoretical results relating to this work are given here; the rest are given in the next chapter, where the approach taken in these two chapter is compared with that taken by other authors.

4.1 Practical Reasoning

In chapter 1, I suggested that there were two types of argument we wanted to generate: those for or against a treatment, and those about the expected effects of a treatment. Many of the rules I have presented have been about some treatment, but I have been vague about where this information comes from. The primary source of information on the effects of treatments comes from the results of clinical trials, published in peer-reviewed journals. These can be formalised as epistemic rules, and describe absolute risk or changes in risk in relation to some treatment, rather than telling us about which treatment to take. Other sources of information, such as guidelines and analyses of past practice (which we can formalise as heuristic rules), might tell us what treatment we should give, but guidelines are often over-simplified and do not incorporate patient values, while analyses of past practice tell us what was done, not what should be done. In theory, guidelines (and practice) should be based on the results of clinical trials, but it would be preferable to reason directly with the results of trials to develop arguments about treatments. However, as we shall see, knowledge in this form requires additional processing before we can use it as we intend.

Example 4.1.1. Consider a statement such as ‘A trial published in the BMJ in 1999 showed that in people with breast cancer, tamoxifen reduces mortality’.

Informally, most of us could use this as the basis of an argument for giving Tamoxifen, because we would like to give (or receive) an intervention that increased overall survival. However, our system, in common with other argumentation frameworks, does not allow us to do this.

Example 4.1.2. If we were to formalise the statement above as a rule, we might write it as:

BMJ1999: People(x)
$\wedge$ hasDisease(x,y) $\wedge$ BreastCancerTypes(y)
$\wedge$ hasTreatment(x,TamE5YrCourse)
$\Rightarrow$ hasDeltaRisk(x,IncreasedOS1.1)

If we used this rule as the support of an argument, it would allow us to generate arguments for IncreasedOS1.1. However, the claim of an argument is always the head of a rule in its support. Therefore we cannot develop arguments for TamE5YrCourse from rules such as this, because it never appears in the head of a rule and so can never appear in the claim of an argument. This is a serious deficiency, and one recognised by other authors ([10], [4] , [8]). A comparison with some of these approaches is given in chapter 5.

Note that the basis for an informal argument to give tamoxifen is dependent upon how we value IncreasedOS1.1; it is only the fact that we desire a state of reduced mortality that makes us argue for tamoxifen. This is the practical syllogism: If X leads to Y, and we desire Y, then we should do X. Such arguments are often referred to as practical arguments, but this term is overloaded, and problematic: a practical argument may be one based on the practical syllogism, or one based on some other practical consideration (limited resources, etc.). For this reason, given that epistemic rules describe beliefs, I use the term ergonic (from $\epsilon\rho\gamma o\nu$ for 'action') to describe rules (and hence arguments) that are about action. We also note a difference in our expectations about the terms People(x) and hasTreatment(x,y). We cannot chose whether People(p) is true, but we can choose whether hasTreatment(p, TamE5YrCourse) holds or not. A more detailed discussion of the practical syllogism is given in [8]. There is also another strand of work, based on the values advanced by certain arguments, where these values may be used to arbitrate between arguments [11] . However, my use of values is distinct from this - a value in OAF is something that a patient, doctor or other interested entity decides is of interest and whose presence or absence is significant. This approach is extended below.

This chapter presents a technique to address the problem of developing ergonic arguments. To do so, I rewrite the epistemic rules into a form which comes closer to capturing the informal argument, and then use the existing definitions for argument construction. This has two advantages: it allows the use of the same mechanisms for argument construction, and it allows us to author a single rulebase, from which we develop two types of rules. The process depends upon two components: Valuation and Intention, and I deal with these in turn.

4.2 Valuation

4.2.1 An informal approach

Let us start by reconsidering the rule I presented earlier:

BMJ1999: People(x)

$\wedge$ hasDisease(x,y) $\wedge$ BreastCancerTypes(y)

$\wedge$ hasTreatment(x,TamE5YrCourse)

$\Rightarrow$ hasDeltaRisk(x,IncreasedOS1.1)

As I suggested above, it is the fact that increased overall survival is of interest that would allow such a rule as the basis for an (informal) argument for tamoxifen. However, this 'interest' has two aspects, relevance and desirability. Consider the two rules given below:

JUS1990: People(x): hasDisease(x,y) $\wedge$ BreastCancerTypes(y)

$\wedge$ hasTreatment(x,z) $\wedge$ Tamoxifen(z)

$\wedge$ DecreasedOS(w)

$\Rightarrow$ hasDeltaRisk(x,w)

JIF2000: People(x): hasDisease(x,y) $\wedge$ BreastCancerTypes(y)

$\wedge$ hasTreatment(x,z) $\wedge$ Tamoxifen(z)

$\wedge$ IncreasedRiskPinkUrine(w)

$\Rightarrow$ hasDeltaRisk(x,w)

The first says that Tamoxifen decreases survival; most would see it as a good reason not to take Tamoxifen, as the outcome would have relevance, but undesirability. The second says that Tamoxifen increases the risk of having pink urine; many would consider it as the basis for neither an argument for or against Tamoxifen, due to its irrelevance. I shall refer to this combination of relevance and desirability as being of 'value', but in our framework we need to consider what relevance and desirability mean in terms of sets of formulae. Although this is useful, it does not help us to decide *whose* values we consider, *what* to value, nor *how* to represent the values. I deal with these below.

4.2.1.1 Whose values ?

There are many sets of values (from patients, doctors, managers, etc.) which may conflict and compete. However, I shall consider a patient's values as being of primary importance. This does not diminish the importance or challenge of resolving these competing viewpoints, but it is not the focus of this thesis. Even patients' values are not uniform - patients differ from each other, and any single patient's values may vary over time. For this reason, we cannot expect to establish a set of universal values, although one might start with a common set (one could imagine an 'NHS default', for example) and then refine them from there.

4.2.1.2 What can we value ?

Our epistemic rules have two different types of heads: for two constants $t_1, t_2 \in I$, either hasRisk(t_1, t_2) or hasDeltaRisk(t_1, t_2).

Example 4.2.1. Consider the following rules.

RoswellPark: Women(p)

$\wedge$ hasDisease(p,q) $\wedge$ BreastCancer(q)

$\wedge$ hasTreatment(p,s) $\wedge$ Tamoxifen(s)

$\wedge$ RiskEndoCa(r) $\wedge$ refersDisease(r,t)

$\wedge$ EndometrialCa(t) $\wedge$ hasValue(r,0.005)

$\Rightarrow$ hasRisk(p,r)

BMJ1999: People(x)

$\wedge$ hasDisease(x,y) $\wedge$ BreastCancerTypes(y)

$\wedge$ hasTreatment(x,TamE5YrCourse)

$\Rightarrow$ hasDeltaRisk(x,IncreasedOS1.1)

Both express a relationship between a set of conditions (in the body) and risk. However, the former describes an absolute risk of an outcome for some group of patients, while the second describes the relative change in risk for some group of patients. The differences between the two types of statement can be understood in different ways, but a simple categorisation might be that the first might be the result of a cohort study that followed a group of patients for some period of time, while the second is the sort of data that would come from a clinical trial, where we introduce an intervention and assess the impact of the intervention of risk. For full discussion of these matters, and the relationship between the different outcomes that can be measured in different types of study, the reader is referred to any standard textbook on medical statistics.

Expressing a preference for an absolute level of risk is problematic, as it is difficult to like or dislike an absolute risk of some outcome. We might say that a risk of death of 10% is higher than we would like, but for someone with a large inflammatory breast cancer, such an outcome would be better than average, whereas for someone with a cough, it would seem extremely high. The problem is that our valuation of an absolute risk level is very context and situation dependent. On the other hand, if we had some intervention that resulted in a 10% *relative* increase in the risk of death, we might generally agree that such an intervention was not desirable, irrespective of the baseline mortality. In the end, we might still accept the intervention because of other over-riding concerns, but the increase in the risk of death would remain a drawback. In the breast cancer ontology, absolute risk relations are represented by hasRisk(t_1, t_2) and change in risk by hasDeltaRisk(t_1, t_2). We may not value all changes in risk (for example, the change in risk of pink urine, above), but everything we value is a change in risk.

4.2.1.3 Representing Values

Now I have defined what we will value, and once we know a patient's values, I need a technique to represent them in OAF. Since all the other elements of the language are in the ontology, it seems natural to enforce this constraint for the patients' values as well. Since we are dealing with individual patients, we want to relate values to patients, and the simplest way to do this is to list a patient's values, as a set of literals in the language.

This captures the relevance of an outcome, but as suggested above, valuation is based on both relevance and desire. We therefore need two lists - one of outcomes that are relevant and desirable and the other of things that are relevant and undesirable. Logical negation does not allow us to combine these in one; consider the fact that while we may not desire Decreased overall survival (as an idea), ¬DecreasedOS (as an ontological term) is the set of things in Δ disjoint with DecreasedOS - including people, drugs and many other things; thus we use two lists: those things that are ***positively valued***, and those that are ***negatively valued***. This is different to some of the previous work, such as [10]. Finally, we should expect our two lists of values to be disjoint. A patient might change their values (thus moving an outcome from one set to another), but at any one time, the two lists should contain separate sets of formulae.

4.2.2 A Formal Approach

As discussed above, our valuation consists of two groups of things, those that we positively and negatively value. These things are expressed as literals in $\mathcal{G}_{\mathcal{K}}$. As discussed above, I shall value only those rules which contain hasDeltaRisk(t_1,t_2) in the head.

Example 4.2.2. For example, MsJones could chose to value "a 10% increase in overall survival" , in which case we might say our values were:

hasDeltaRisk(MsJones,IncreasedOS1.1)

Earlier, I said that the preferences that we express reflect that preferences of individual patients. In order to be able to describe this, we will need to be able to refer to individual patients in our ontology and hence OAF. I refer to patients by their name.

Definition 4.2.3. Let $\mathcal{K}$ be some vocabulary, and let I be the set of concrete individuals in $\mathcal{K}$. Let Ω = (Θ, Δ) be a S-OAF. The function $PeopleNames : \Omega \mapsto I$ returns the set of people in an OAF. That is, for all $t \in I$ such that $\Delta \vdash_{Ont}$ People(t), $t \in PeopleNames(\Omega)$.

Since this thesis is largely concerned with the effect of treatment on outcomes, it will be useful to have a similar function to refer to treatments.

Definition 4.2.4. Let $\mathcal{K}$ be some vocabulary, and let I be the set of concrete individuals in $\mathcal{K}$. Let Ω = (Θ, Δ) be a S-OAF. The function $Treatments : \Omega \mapsto I$ returns the set of treatments in an OAF. That is, for all $t \in I$ such that $\Delta \vdash_{Ont}$ TreatmentTypes(t), $t \in Treatments(\Omega)$.

The same is true of changes in risk.

Definition 4.2.5. Let $\mathcal{K}$ be some vocabulary, and let I be the set of concrete individuals in $\mathcal{K}$. Let Ω = (Θ, Δ) be a S-OAF. The function $DeltaRisks : \Omega \mapsto I$ returns the set of changes in risk in an OAF. That is, for all $t \in I$ such that $\Delta \vdash_{Ont}$ DeltaRisks(t), $t \in DeltaRisks(\Omega)$.

We can now refer to people by name and so associate values with individuals. For some patient p, $\mathcal{V}_p^+$ is the set of positive values for patient p, $\mathcal{V}_p^-$ is the set of negative values for patient p and V_p = $\mathcal{V}_p^+$ $\cup$ $\mathcal{V}_p^-$ is the set of values for patient p.

Definition 4.2.6. Let $\mathcal{K}$ be some vocabulary and $\mathcal{G}_{\mathcal{K}}$ the set of literals. Let Ω be a S-OAF. Then for a patient, p, s.t. $p \in PeopleNames(\Omega)$, let $\mathcal{V}_p^+$, $\mathcal{V}_p^-$ be sets of literals in $\mathcal{G}_{\mathcal{K}}$ such that $\mathcal{V}_p^+$, $\mathcal{V}_p^- \subseteq DeltaRisks(\Omega)$ and $\mathcal{V}_p^+ \cap \mathcal{V}_p^- = \emptyset$.

Example 4.2.7. Consider a patient, Ms. Jones. When asked, she specifies her positive values as "increased overall survival" and "increased breast cancer disease-free survival". In Δ_{BC}, these are both classes in the ontology, and there are many different instances of each of these classes, which taken together form the formal representation of her positive values. Ms. Jones' set of positive intentions ($\mathcal{V}^+_{MsJones}$) would then be:

$\mathcal{V}^+_{MsJones}$ ={hasDeltaRisk(MsJones,IncreasedOS1.01),

hasDeltaRisk(MsJones,IncreasedOS1.02),

hasDeltaRisk(MsJones,IncreasedOS1.03),

hasDeltaRisk(MsJones,IncreasedOS1.1),

hasDeltaRisk(MsJones,IncreasedOS1.2),

hasDeltaRisk(MsJones,IncreasedBrCaDFS1.05),

hasDeltaRisk(MsJones,IncreasedBrCaDFS1.1),

hasDeltaRisk(MsJones,IncreasedBrCaDFS1.2),

hasDeltaRisk(MsJones,IncreasedBrCaDFS1.21),

hasDeltaRisk(MsJones,IncreasedBrCaDFS1.3),

hasDeltaRisk(MsJones,IncreasedBrCaDFS1.5),

hasDeltaRisk(MsJones,IncreasedBrCaDFS1.7)}.

Similarly, we can give a set of negative values.

Example 4.2.8. Ms. Jones decides that she negatively values an "increased risk of endometrial cancer", or "decreased overall survival" . Her set of negative values ($\mathcal{V}^-_{MsJones}$) will then be:

$\mathcal{V}^-_{MsJones}$ ={hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer2.3),

hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer2.53),

hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer4.1),

hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer6.0),

hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer6.4),

hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer7.5)

hasDeltaRisk(MsJones,DecreasedOS0.7)}.

I started this chapter with a discussion of how informal arguments depend on the valuation of the expected outcome, and have developed a notation to record a patient's values, and discussed what we can value. The next section discusses the second aspect, Intention.

4.3 Intention

4.3.1 Databases and Intentions

As I mentioned at the beginning of the last chapter, the ontology provides a source of facts on which we can base our arguments. It is useful for recording a set of facts about various patients (e.g. Ms. Jones), and also for arguing about (for example) what type of cancer her breast lump might be. It is when we come to talk about things that we intend to do, and so are are not yet true, that the simple A-Box representation becomes more difficult.

This may be because of the lack of temporal information in our ontology, and if we included suitable definitions for temporal events and periods, we could perhaps solve some of these problems. For example, we could describe a treatment course as starting at some point in the future in order to capture the idea that it is not yet true, but that we are interested in doing it. Unfortunately, there are three problems with this approach. The first is that if we marked 'potentiality' simply by date of onset, we would need to include in all the rules a 'guard' condition that restricted them to only considering the effects of past and current treatments. The second is that there is a distinction between things we desire, and things in the future: not everything in the future is something we desire, and not everything we desire is in the future. The third is that the 'future' events may prompt preparatory actions, and so a course of action that we were considering might prompt a set of subsidiary actions; furthermore, if we decided not to do that action, we would then have to 'undo' these subsidiary actions. In all three cases, it seems that simple temporal distinctions do not capture what we want.

4.3.2 Belief, Desires, Intentions

The fact that even the addition of a temporal aspect to the ontology would not resolve this problem of the division between things that are true and possible courses of action should lead us to some consideration of the cause. For the last 20 years, psychologists, logicians and computer scientists have, from varying directions, been working on the Belief, Desires and Intention (BDI) model of agent behaviour and it has, at least in general terms, become well-known. It has a variety of different interpretations depending on the school of the different developers, and so is often characterised as either a true representation of human cognitive states [16], an inspiration for a logic [30], or a convenient shorthand for programmers [43]. It is not my intention to give a detailed breakdown of the different approaches to the BDI models, nor their formal models nor their implementations. However, in general the BDI model divides agent cognition into three states: Informational state (Beliefs), Motivational state (Desires) and Deliberative state (Intentions). The precise meanings of each of these categories has been extensively debated (for example, the difference between 'Goals' and 'Desires'), and refinements have been offered [31] but in general the categorisation has proved remarkably stable and has been the basis for several different implemented systems. I would suggest that the reason for the difficulties outlined above is that I am describing something very similar. Indeed, the stability, and utility, of the BDI approach suggests that we could use it as 'frame' in which to hang OAF. When we do so, it fits quite well. The beliefs equate to facts, while desires are mapped to our values. However, OAF has nothing that fits with Intention. If the BDI model is useful, and if OAF is treading on similar ground, then the development of an explicit category of Intention in OAF may be sensible. However, I intend to do this in a manner consistent with the existing presentation of OAF. What this means in practice is that the terms of our intentions must, like the other terms in OAF, be describable within the ontology.

4.3.3 Ontology & Intention

The simplest approach might seem to be to develop some class of instances, Intention(t_1), to denote a patient's intention for some course of action t_1 (where $t_1 \in I$). However, this leads to the same problems as we saw above.

Example 4.3.1. If we say Ms. Jones is taking 2 years of tamoxifen, then we use syntax of the form:

People(MsJones) $\wedge$ hasTreatment(MsJones,TamD2Yr)

which seems clear enough. However, if Intention() is a class, then if Ms. Jones develops an intention to take tamoxifen, we would use a syntax:

People(MsJones) $\wedge$ hasTreatment(MsJones,TamD2Yr) $\wedge$ Intention(TamD2Yr)

This approach has two drawbacks. Firstly Ms. Jones still appears to be the same as someone who is taking Tamoxifen, and so will satisfy rules written for those who are currently taking Tamoxifen. Secondly it involves us making an assertion that TamD2Yr is a type of intention, which is not how it is currently defined in the ontology. The solution is to use a different property to link people and treatment and so signify the intention to take a treatment, as this example shows:

Example 4.3.2. In order to express that she has the intention to take a treatment of 2 years of Tamoxifen, I use the formula

People(MsJones) $\wedge$ hasPosIntent(MsJones, TamD2Yr)

and in order to express the fact that she has the intention *not* to take a treatment of the type tamoxifen, I use the formula

People(MsJones) $\wedge$ hasNegIntent(MsJones, TamD2Yr)

Based on this approach, I can give a definition of intention. Many clinical databases might not contain the necessary terminology; this is the reason that I insisted on the DL Roles hasPosIntent(x,y) and hasNegIntent(x,y) being in our ontology, Δ, in Chapter 2. Once we know that Δ contains such terms, we can also be sure that $\mathcal{G}_{\mathcal{K}}$ contain the necessary predicates.

Definition 4.3.3. Let $\mathcal{K}$ be some vocabulary, $\mathcal{G}_{\mathcal{K}}$ be the set of literals and I the set of concrete individuals in $\mathcal{K}$, with $t_1, t_2 \in I$. A ***positive intention*** is a literal $\phi_1 \in \mathcal{G}_{\mathcal{K}}$ of the form:

ϕ_1 = hasPosIntent(t_1, t_2)

where $t_1 \in PeopleNames(\Omega)$ and $t_2 \in TreatmentTypes(\Omega)$ both hold

Example 4.3.4. Let ϕ_1 = hasPosIntent(MsJones,TamD2YrCourse). Then ϕ_1 is a positive intention.

Definition 4.3.5. Let $\mathcal{K}$ be some vocabulary, $\mathcal{G}_{\mathcal{K}}$ be the set of literals and I the set of concrete individuals in $\mathcal{K}$, with $t_1, t_2 \in I$. A ***negative intention*** is a literal $\phi_2 \in \mathcal{F}_{\mathcal{K}}$ of the form

ϕ = hasNegIntent(t_1, t_2)

where $t_1 \in PeopleNames(\Omega)$ and $t_2 \in TreatmentTypes(\Omega)$ both hold

Example 4.3.6. Let ϕ_2 = hasNegIntent(MsJones,TamFgt5Yr). Then ϕ_2 is a negative intention

Although I have shown how we can use two predicates to describe intentions, these predicates are still in the ontology, and so can be translated into DL formula. Not only is this in keeping with the 'spirit' of OAF, it also means that we can use our ontology based reasoning (for example, conflict detection and entailment) with intentions as with any other formulae. Since I want to retain a close mapping between beliefs and intentions (as we shall see in the next chapter), using the same format has its advantages. It also points to the link between intention and actions, especially for long-lasting actions (such as a course of drugs). For some things in the A-Box, we accept the fact that they happened, without our agency or desire (such as developing breast cancer). However, for those things that we can decide on, it should be a minimal criterion that at least at some point, we developed an intention for them. Indeed, we can use the development of intentions as a check on our current and past performance by asking to what extent do the treatments in the A-Box correspond with what we would now suggest? It is this sort of use of intention, primarily as a marker for intent, rather than just 'futureness', that requires more than a simple temporal encoding.

4.4 Valuation & Intention for End-Means Valuation

So far this chapter has described an approach to valuation and intention. This section uses both of these to formalise the practical syllogism. To do this, we need a valuation function which decides whether a literal is valued. Above, I defined two sets of values, $\mathcal{V}_P^+$ and $\mathcal{V}_P^-$, to represent patient values. I therefore use these sets to decide the valuation status of some formula.

Definition 4.4.1. Let $\mathcal{K}$ be some vocabulary, $\mathcal{G}_{\mathcal{K}}$ the set of literals and Ω be a S-OAF. Then for a patient, p, s.t. $p \in PeopleNames(\Omega)$, let V_p be the patient's values and $\phi \in \mathcal{G}_{\mathcal{K}}$ be a literal. Then the function $Valuation$: $\mathcal{G}_{\mathcal{K}} \mapsto \{Neg, Pos, Unvalued\}$ provides the ***valuation*** of ϕ accordingly:

If $\phi \in \mathcal{V}_p^+$ then $Valuation(\phi) = Pos$

else if

$\phi \in \mathcal{V}_p^-$ then $Valuation(\phi) = Neg$

Otherwise, $Valuation(\phi) = Unvalued$

Example 4.4.2. Considering Ms. Jones again. If, as above,

$\mathcal{V}_{MsJones}^+$ ={hasDeltaRisk(MsJones,IncreasedOS1.01),

hasDeltaRisk(MsJones,IncreasedOS1.02),

hasDeltaRisk(MsJones,IncreasedOS1.03),

hasDeltaRisk(MsJones,IncreasedOS1.1),

hasDeltaRisk(MsJones,IncreasedOS1.2),

hasDeltaRisk(MsJones,IncreasedBrCaDFS1.05),

hasDeltaRisk(MsJones,IncreasedBrCaDFS1.1),

hasDeltaRisk(MsJones,IncreasedBrCaDFS1.2),

hasDeltaRisk(MsJones,IncreasedBrCaDFS1.21),

hasDeltaRisk(MsJones,IncreasedBrCaDFS1.3),

hasDeltaRisk(MsJones,IncreasedBrCaDFS1.5),

hasDeltaRisk(MsJones,IncreasedBrCaDFS1.7)}.

and

$\mathcal{V}^{-}_{MsJones}$ ={hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer2.3),

hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer2.53),

hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer4.1),

hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer6.0),

hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer6.4),

hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer7.5)

hasDeltaRisk(MsJones, DecreasedOS0.7)}.

and we want to determine the valuation of a formula ϕ_3 such that:

ϕ_3 = hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer4.1)

Then $\phi_3 \in \mathcal{V}^{-}_{MsJones}$, and thus

$Valuation$(hasDeltaRisk(MsJones,IncreasedRiskEndometrialCancer4.1), V_p) = Neg

We now have definitions of patient values and intentions, and a function to determine whether a ground formula is positively or negatively valued. These are used to rewrite epistemic rules.

4.4.1 Rewriting Rules

I now describe how we develop intentions from a combination of existing rules and values. I start with the rule syntax:

$$Label : Body \Rightarrow Head$$

From our earlier definition, $Body$ is some (possibly null) formula. However, as suggested above the predicates in the body of the rule can be divided into those about which we have no choice (such as People()) and those about which we have a choice (such as Tamoxifen()). This division is not formally defined, but the point is that we divide the predicates into those that we consider to be optional and those that we don't. This division may change from domain to domain, but the idea of differentiating 'choosable' from 'non-choosable' things remains, and in our domain, only types of treatment will be considered to be choosable. Recall that the definition of an epistemic rule in chapter 3 limited such rules to having at most one literal of the form TreatmentTypes() in the body of the rule, and we can see that it is simple to transform an epistemic rule into an ergonic one, as the example below demonstrates:

Example 4.4.3. Let $\mathcal{K}$ be a vocabulary and $\mathcal{G}_{\mathcal{K}}$ the set of literals. Then $\Gamma \subseteq \mathcal{G}_{\mathcal{K}}$ is a set of facts, and Δ_{BC} is the breast cancer ontology. Then Ω =(Θ,$\Delta \cup \Gamma$) is a S-OAF, where Γ and Θ are as below:

Γ = {Woman(MsJones), hasDisease(MsJones,ProtoBreastCancer)}

Θ= {BMJ1999}

where

BMJ1999: People(x): hasDisease(x,y) $\wedge$ BreastCancerTypes(y)
$\wedge$ hasTreatment(x,z) $\wedge$ Tamoxifen(z)
$\wedge$ IncreasedOS(w)
$\Rightarrow$ hasDeltaRisk(x,w)

Given a set of values such that:

$\mathcal{V}^{+}_{MsJones}$ ={hasDeltaRisk(MsJones,IncreasedOS1.01),
hasDeltaRisk(MsJones,IncreasedOS1.02),
hasDeltaRisk(MsJones,IncreasedOS1.03),
hasDeltaRisk(MsJones,IncreasedOS1.1),
hasDeltaRisk(MsJones,IncreasedOS1.2)}

Then we could rewrite our rule, BMJ1999, based on Ms. Jones' values, so that it becomes a rule for the intervention:

BMJ1999*: People(x)
$\wedge$ hasDisease(x,y) $\wedge$ BreastCancerTypes(y)
$\wedge$ Tamoxifen(z)
$\Rightarrow$ hasPosIntent(x,z)

As is hopefully clear, the transformation from one type of rule to the other is relatively simple, and so I now (finally) present a technique to rewrite rules to handle values and intentions, which results in the rewriting of epistemic rules as ergonic ones. To do so I define a value-based re-writing function, $Rewrite()$ that takes an epistemic rule, and produces an ergonic one. Note that there are four cases covered by the function, two for positively and negatively valued rules, and two for rules which either include a hasValue() predicate in the body or do not. I start by considering rules whose bodies contain hasValue(), recalling that hasValue() refers to the value associated with an outcome, not a patient's values. Referring to the numbered points below, we ensure that the rule is of the correct form (1), check that the individuals are of the correct type (2,3) and then check that the body of the rule includes a treatment (4), and a change in risk (5), which has an associated numeric value (6). If these hold, and we positively value the head of the rule, then we generate a new rule with a head that contains the hasPosIntent(t_1, t_3) predicate, and which lacks the hasTreatment(t_1, t_3), hasValue(t_2, d_1) and DeltaRisk(t_3) predicates in the body. The next case proceeds in an identical fashion, except that we deal with rules whose heads are negatively valued. The remaining two cases deal with rules of similar forms, but which lack the hasValue(t_2, d_1) predicate in the body.

Definition 4.4.4. Let $\mathcal{K}$ be a DL ontology language, I the set of concrete individuals in $\mathcal{K}$ and D the set of datatypes in $\mathcal{K}$, such that $t_{1,}t_{2,},t_3 \in I$, $d_1 \in D$ and $t_1 \in PeopleNames(\Omega)$, $t_2 \in DeltaRisks(\Omega)$ and $t_3 \in Treatments(\Omega)$. Let V_p be a set of patient's values, and Δ be a DL ontology, $\mathcal{R}_{\mathcal{K}}^{\mathcal{P}}$ the set of epistemic rules and $\mathcal{R}_{\mathcal{K}}^{\mathcal{R}}$ the set of all ergonic rules. The ***value-based rewriting function***, $Rewrite$: $\mathcal{R}_{\mathcal{K}}^{\mathcal{P}} \mapsto \mathcal{R}_{\mathcal{K}}^{\mathcal{R}}$ is defined as follows:

Case 1: For a rule $r \in \mathcal{R}_{\mathcal{K}}^{\mathcal{P}}$, such that

1. $Head(r)$ = hasDeltaRisk(t_1,t_2)
2. hasTreatment(t_1, t_3) $\in Body(r)$
3. For some class $A \in C$, $A(t_2) \in Body(r)$
4. $d_1 \in D$ and hasValue(t_2, d_1) $\in Body(r)$

then if $Valuation(Head(r), V_p) = Pos$ then $Rewrite(r)$ is of the form

$Label(r)^*$:

$\bigwedge(((Body(r))\backslash$ {hasTreatment(t_1, t_3), $A(t_2)$, hasValue(t_2, d_1)})

$\Rightarrow$ hasPosIntent(t_1,t_3)

Case 2: For a rule $r \in \mathcal{R}_{\mathcal{K}}^{\mathcal{P}}$, such that

1. $Head(r)$ = hasDeltaRisk(t_1,t_2)
2. hasTreatment(t_1, t_3) $\in Body(r)$
3. For some class $A \in C$, $A(t_2) \in Body(r)$
4. $d_1 \in D$ and hasValue(t_2, d_1) $\in Body(r)$

then if $Valuation(Head(r), V_p) = Neg$ then $Rewrite(r)$ is of the form

$Label(r)^*$:

$\bigwedge(((Body(r))\backslash$ {hasTreatment(t_1, t_3), $A(t_2)$, hasValue(t_2, d_1)})

$\Rightarrow$ hasNegIntent(t_1,t_3)

Case 3: For a rule $r \in \mathcal{R}_{\mathcal{K}}^{\mathcal{P}}$, such that

1. $Head(r)$ = hasDeltaRisk(t_1,t_2)

2. For some class $A \in C$, $A(t_2) \in Body(r)$

3. hasTreatment(t_1, t_3) $\in Body(r)$

then if $Valuation(Head(r), V_p) = Pos$ then let $Rewrite(r)$ is of the form

$Label(r)^*$:

$\bigwedge(((Body(r))\setminus$ {hasTreatment(t_1, t_3), $A(t_2)$})

$\Rightarrow$ hasPosIntent(t_1,t_3)

Case 4: For some $r \in \mathcal{R}^{\mathcal{P}}_{\mathcal{K}}$, such that

1. $Head(r)$ = hasDeltaRisk(t_1,t_2)

2. For some class $A \in C$, $A(t_2) \in Body(r)$

3. hasTreatment(t_1, t_3) $\in Body(r)$

then if $Valuation(Head(r), V_p) = Neg$ then $Rewrite(r)$ is of the form

$Label(r)^*$:

$\bigwedge(((Body(r))\setminus$ {hasTreatment(t_1, t_3), $A(t_s)$})

$\Rightarrow$ hasNegIntent(t_1,t_3)

Note that for a rule r where none of case 1 - 4 hold, $Rewrite(r)$ returns $null$.

From this point forward, I will begin to make the examples increasingly realistic. In particular, I will start using many of the terms that occur in the case study in chapter 7, so as to make them more familiar. An example of this occurs below, where I talk about TamoxifenByLength. This is because trials of tamoxifen involved both variations in dose and variations in length of treatment, and the rules below apply only to different durations of treatment with tamoxifen, not variations in dose.

Example 4.4.5. Let $\mathcal{K}$ be some vocabulary, $\mathcal{G}_{\mathcal{K}}$ the set of literals, Γ, $\mathcal{V}^{+}_{MsJones}$, $\mathcal{V}^{-}_{MsJones} \subset \mathcal{G}_{\mathcal{K}}$ be sets of literals and Δ_{BC} the breast cancer ontology. Let Ω = $(\Theta, \Delta_{BC} \cup \Gamma)$ be a S-OAF such that

Γ = $\{$Woman(MsJones), hasDisease(MsJones,ProtoBreastCancer)$\}$

Θ= $\{$BMJ1999, NEJM2001$\}$

where

BMJ1999: People(x)

$\wedge$ hasDisease(x,y) $\wedge$ BreastCancerTypes(y)

$\wedge$ hasTreatment(x,z) $\wedge$ Tamoxifen(z)

$\wedge$ IncreasedOS(w)

$\Rightarrow$ hasDeltaRisk(x,w)

NEJM2001: People(x)

$\wedge$ hasTreatment(x,y) $\wedge$ Tamoxifen(y)

$\wedge$ IncreasedRiskEndometrialCa(z)

$\Rightarrow$ hasDeltaRisk(x,z)

and

$\mathcal{V}^{+}_{MsJones}$ ={hasDeltaRisk(MsJones,IncreasedOS1.01),

hasDeltaRisk(MsJones,IncreasedOS1.02),

hasDeltaRisk(MsJones,IncreasedOS1.03),

hasDeltaRisk(MsJones,IncreasedOS1.1),

hasDeltaRisk(MsJones,IncreasedOS1.2)}

and

$\mathcal{V}^{-}_{MsJones}$ ={hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer2.3),

hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer2.53),

hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer4.1),

hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer6.0),

hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer6.4),

hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer7.5)}.

Then if we apply $Rewrite()$ to the elements of Θ in turn, we get:

BMJ1999*: People(x)

$\wedge$ hasDisease(x,y) $\wedge$ BreastCancerTypes(y)

$\wedge$ Tamoxifen(z)

$\Rightarrow$ hasPosIntent(x,z)

NEJM2001*: People(x)

$\wedge$ TamoxifenTypes(y)

$\Rightarrow$ hasNegIntent(x,y)

So far, we have regarded a patients values as being 'external' to a simple OAF. In some circumstances, we shall want to consider them as part of an OAF, and this type of OAF is called an 'extended OAF'

Definition 4.4.6. Let $\mathcal{K}$ be a vocabulary, $\mathcal{O}_{\mathcal{K}}$ the set of ontology formulae, $\mathcal{R}_{\mathcal{K}}$ the set of defeasible rules, $\mathcal{G}_{\mathcal{K}}$ the set of literals, and $\Omega = \{\Theta, \Delta\}$ be a S-OAF. Also, let p be in $PeopleNames(\Omega)$, $\mathcal{V}^{+}_{P} \subseteq \mathcal{G}_{\mathcal{K}}$ be the set of positive values for patient p and $\mathcal{V}^{-}_{p} \subseteq \mathcal{G}_{\mathcal{K}}$ the set of negative values for patient p. An ***Extended Ontology-based Argumentation Framework*** (E-OAF), denoted Ω, is a tuple (Θ, Δ, V_p) where $\Theta \subset \mathcal{R}_{\mathcal{K}}$ is a finite set of defeasible rules, $\Delta \subseteq \mathcal{O}_{\mathcal{K}}$ is an ontology and V_p = ($\mathcal{V}^{+}_{p}$, $\mathcal{V}^{-}_{p}$).

One of the advantages of the use of the $Rewrite$ function is that we can author a single set of epistemic defeasible rules, and then develop a set of ergonic rules depending on the patient's values. This is shown in the example below where Θ is a set of *epistemic* defeasible rules.

Example 4.4.7. Let $\mathcal{K}$ be a vocabulary, $\mathcal{G}_{\mathcal{K}}$ the set of literals, $\Gamma \subseteq \mathcal{G}_{\mathcal{K}}$ a set of facts, Θ a set of epistemic defeasible rules and Δ_{BC} the breast cancer ontology. Let $V_{MsJones}$ be the sets of positive and negative values for MsJones, and Φ be set of ergonic defeasible rules such that $\Phi = \{Rewrite(r|r \in \Theta$ and $Rewrite(r){\neq}null)$ }. Then Ω = ($\Theta \cup \Phi$, $\Delta_{BC} \cup \Gamma$, V_p) is an E-OAF, where Θ, Φ, Γ and $V_{MsJones}$ are as below.

Γ = {Woman(MsJones), hasDisease(MsJones,ProtoBreastCancer)}

Θ= {BMJ1999, NEJM2001}

BMJ1999: People(x): hasDisease(x,y) $\wedge$ BreastCancerTypes(y)
$\wedge$ hasTreatment(x,z) $\wedge$ TamoxifenByLength(z)
$\wedge$ IncreasedOS(w)
$\Rightarrow$ hasOutcome(x,w)

NEJM2001: People(x)
$\wedge$ hasTreatment(x,y) $\wedge$ TamoxifenByLength(y)
$\wedge$ IncreasedRiskEndometrialCa(z)
$\Rightarrow$ hasDeltaRisk(x,z)

$\mathcal{V}^{+}_{MsJones}$ ={hasDeltaRisk(MsJones,IncreasedOS1.01),
hasDeltaRisk(MsJones,IncreasedOS1.02),
hasDeltaRisk(MsJones,IncreasedOS1.03),
hasDeltaRisk(MsJones,IncreasedOS1.1),
hasDeltaRisk(MsJones,IncreasedOS1.2)}

and

$\mathcal{V}^{-}_{MsJones}$ ={hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer2.3),
hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer2.53),
hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer4.1),
hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer6.0),
hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer6.4),
hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer7.5)
hasDeltaRisk(MsJones, DecreasedOS0.7)}.

$V_{MsJones}$ = ($\mathcal{V}^{+}_{MsJones}$, $\mathcal{V}^{-}_{MsJones}$)

The application of $Rewrite()$ to the elements of Θ gives us:

Φ = {BMJ1999*, NEJM2001*}

BMJ1999*: People(x)

$\wedge$ hasDisease(x,y) $\wedge$ BreastCancerTypes(y)

$\wedge$ TamoxifenByLength(z)

$\Rightarrow$ hasPosIntent(x,z)

NEJM2001*: People(x)

$\wedge$ TamoxifenByLength(y)

$\Rightarrow$ hasNegIntent(x,y)

In developing the arguments we note that the bodies of BMJ1999 and NEJM2001 are not satisfied by $(\Delta_{BC} \cup \Gamma)$. Therefore we need concern ourselves only with ergonic rules. In the interests of clarity, these are given below:

BMJ1999$_1{}^{*1}$:People(MsJones)

$\wedge$ hasDisease(MsJones,ProtoAdenoCaBreast) $\wedge$ BreastCancerTypes(ProtoAdenoCaBreast)

$\wedge$ TamoxifenByLength(TamD2YrCourse)

$\Rightarrow$ hasPosIntent(MsJones,TamD2YrCourse)

BMJ1999$_2{}^{*2}$:People(MsJones)

$\wedge$ hasDisease(MsJones,ProtoAdenoCaBreast) $\wedge$ BreastCancerTypes(ProtoAdenoCaBreast)

$\wedge$ TamoxifenByLength(TamE5YrCourse)

$\Rightarrow$ hasPosIntent(MsJones,TamE5YrCourse)

BMJ1999$_3{}^{*3}$:People(MsJones)

$\wedge$ hasDisease(MsJones,ProtoAdenoCaBreast) $\wedge$ BreastCancerTypes(ProtoAdenoCaBreast)

$\wedge$ TamoxifenByLength(TamFgt5YrCourse)

$\Rightarrow$ hasPosIntent(MsJones,TamFgt5YrCourse)

NEJM2001$_1{}^{*1}$:(People(MsJones)

$\wedge$ EndometrialCaTypes(ProtoEndometrialCa)

$\wedge$ TamoxifenByLength(TamD2YrCourse)

$\Rightarrow$ hasNegIntent(MsJones,TamD2YrCourse)

NEJM2001$_2{}^{*2}$:(People(MsJones)

$\wedge$ EndometrialCaTypes(ProtoEndometrialCa)

$\wedge$ TamoxifenByLength(TamE5YrCourse)

$\Rightarrow$ hasNegIntent(MsJones,TamE5YrCourse)

NEJM2001$_3{}^{*3}$:(People(MsJones)

$\wedge$ EndometrialCaTypes(ProtoEndometrialCa)

$\wedge$ TamoxifenByLength(TamFgt5YrCourse)

$\Rightarrow$ hasNegIntent(MsJones,TamFgt5YrCourse)

Then we are able to construct the arguments:

a_1 $\langle${BMJ1999^{*1}}, hasPosIntent(MsJones,TamD2YrCourse)$\rangle$

a_2 $\langle${BMJ1999^{*2}}, hasPosIntent(MsJones,TamE5YrCourse)$\rangle$

a_3 $\langle${BMJ1999^{*3}}, hasPosIntent(MsJones,TamFgt5YrCourse)$\rangle$

a_4 $\langle${NEJM2001^{*1}}, hasNegIntent(MsJones,TamD2YrCourse)$\rangle$

a_5 $\langle${NEJM2001^{*2}}, hasNegIntent(MsJones,TamE5YrCourse)$\rangle$

a_6 $\langle${NEJM2001^{*3}}, hasNegIntent(MsJones,TamFgt5YrCourse)$\rangle$

Note that here a_1- a_3 conflict with a_4 - a_6, as well as with each other

This completes the introduction of value-based reasoning. We can see the effect of this in two ways. The first is to give a worked example to show that it answers at least some of our needs. The second is to give some simple proofs, some of which are given here and some at the end of the next chapter. I start by developing the worked example; it is used in the next chapter as well, and so I develop it in some detail.

4.5 Worked Example

Let $\mathcal{K}$ be a vocabulary, $\mathcal{G}_{\mathcal{K}}$ the set of literals, $\Gamma \subseteq \mathcal{G}_{\mathcal{K}}$ a set of facts, Θ a set of epistemic defeasible rules and Δ_{BC} the breast cancer ontology. Let $V_{MsJones}$ be the sets of positive and negative values for MsJones, and Φ be set of ergonic defeasible rules. Then Ω = ($\Theta \cup \Phi$, $\Delta_{BC} \cup \Gamma$, V_p) is an E-OAF, where Θ, Φ, Γ and $V_{MsJones}$ are as below.

Γ = {Woman(MsJones), hasDisease(MsJones,ProtoBreastCancer)}

$\mathcal{V}^+_{MsJones}$ ={hasDeltaRisk(MsJones,IncreasedOS1.01),

hasDeltaRisk(MsJones,IncreasedOS1.02),

hasDeltaRisk(MsJones,IncreasedOS1.03),

hasDeltaRisk(MsJones,IncreasedOS1.1),

hasDeltaRisk(MsJones,IncreasedOS1.2)}

and

$\mathcal{V}^-_{MsJones}$ ={hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer2.3),

hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer2.53),

hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer4.1),

hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer6.0),

hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer6.4),

hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer7.5)

hasDeltaRisk(MsJones, DecreasedOS0.7)}.

$V_{MsJones}$ = ($\mathcal{V}^+_{MsJones}$, $\mathcal{V}^-_{MsJones}$)

Θ = {NICE2003, NLCN1997, NCRN2001, BMJ1999, BJC2004} where

NICE2003: People(x)

∧ hasDisease(x,y) ∧ BreastCancerTypes(y)

∧ Tamoxifen5YrCourse(z)

⇒ hasTreatment(x,z)

NLCN1997: People(x)

∧ hasDisease(x,y) ∧ BreastCancerTypes(y)

∧ Tamoxifen2YrCourse(z)

⇒ hasTreatment(x,z)

NCRN2001: People(x)

∧ hasDisease(x,y) ∧ BreastCancerTypes(y)

∧ ECMFChemotherapy(z)

⇒ hasTreatment(x,z)

BMJ1999: People(x)

∧ hasTreatment(x,y) ∧ TamoxifenByLength(y)

∧ IncreasedOS(z) ∧ hasValue(z,1.2)

⇒ hasDeltaRisk(x,z)

BJC2004: (People(x)

∧ hasTreatment(x,y) ∧ TamoxifenByLength(y)

∧ IncreasedRiskEndometrialCancer(z) ∧ hasValue(z,4.1)

⇒ hasDeltaRisk(x,z)

and

Φ = $\{Rewrite(r|r \in \Theta$ and $Rewrite(r) \neq null)$ }

At this point, let us look at our presentation of Ω. As before, it consists of defeasible rules and facts, associated with an ontology, but now includes a set of ergonic rules and a patient's values. As before, it is based on the breast cancer ontology. However, the rules are different: NICE2003, NLCN1997 and NCRN2001 are heuristic defeasible rules. In chapter three, I suggested that such rules might be generated by modelling guidelines or a review of patient records (which tend to make statements about what should be done) rather than the trial literature directly (which tend to report the effect of doing something). BMJ1999 and BJC2004, however are closer to those developed from journal papers. Although there are potential problems in writing different sort of knowledge in the same rule-base (for example, some of NICE2003, NLCN1997 or NCRN2001 could be based on BMJ1999 or BJC2004) for now I will accept it.

We form the following arguments from Ω:

a_1 $\langle$\{NICE2003\}, hasTreatment(MsJones,TamE5YrCourse)$\rangle$

a_2 $\langle$\{NLCN1997\}, hasTreatment(MsJones,TamD2YrCourse)$\rangle$

a_3 $\langle$\{NCRN2001\}, hasTreatment(MsJones,ECMFTypeInstance)$\rangle$

a_4 $\langle$\{BMJ1999^{*1}\}, hasPosIntent(MsJones,TamD2YrCourse)$\rangle$

a_5 $\langle$\{BMJ1999^{*2}\}, hasPosIntent(MsJones,TamE5YrCourse)$\rangle$

a_6 $\langle$\{BMJ1999^{*3}\}, hasPosIntent(MsJones,TamFgt5YrCourse)$\rangle$

a_7 $\langle$\{BJC2004^{*1}\}, hasNegIntent(MsJones,TamD2YrCourse)$\rangle$

a_8 $\langle$\{BJC2004^{*2}\}, hasNegIntent(MsJones,TamE5YrCourse)$\rangle$

a_9 $\langle$\{BJC2004^{*3}\}, hasNegIntent(MsJones,TamFgt5YrCourse)$\rangle$

On the basis of these arguments, we can then make some new arguments:

a_{10} $\langle$\{NICE2003, BMJ1999\}, hasDeltaRisk(MsJones, IncreasedOS1.2)$\rangle$

a_{11} $\langle$\{NLCN1997, BMJ1999\}, hasDeltaRisk(MsJones, IncreasedOS1.2)$\rangle$

a_{12} $\langle$\{NICE2003, BJC2004\}, hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer4.1)$\rangle$

a_{13} $\langle$\{NLCN1997, BJC2004\}, hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer4.1)$\rangle$

4.6 Matching Our Intuitions

I want to start by considering whether or not the application of valuation to the rules does what we feel it should. As I described in 1.4, there are two main sorts of arguments that we wish to see when choosing treatments:

1. Reasons for (and against) a treatment
2. The expected effects of a treatment

My intention in developing the valuation was to allow rules about (2) to be turned into rules about (1), so as to allow us to argue about treatments. The question, therefore, is has it done so?

The simple answer is 'yes'. As we can see from the example above, a_6 and a_7 are arguments for and against the intention to give tamoxifen. Since the machinery for creating such arguments is relatively simple, it should be possible to create such arguments for other treatments, such as chemo- or radiotherapy. In that sense, therefore, we have satisfied (1), although we will return to some of the details later.

However, there is a problem when we come to (2). As we can see from our example, we have arguments for and against the intention to give tamoxifen. As a clinician I would like to see the impact that giving tamoxifen might have, and fortunately we have rules that tell us this: BMJ1999 and BJC2004 both detail the effect that giving tamoxifen might have (although they disagree on it). Unfortunately, however, this doesn't work - by describing the treatment (in this case tamoxifen) as an intention, it means that it no longer satisfies the conditions in the body of BMJ1999 and BJC2004, and so we never develop arguments about the effects of giving tamoxifen. To do this, we need to (at least temporarily) commit to a course of action, and then examine its consequences. This is developed in the next chapter, and will involve a further extension to the framework.

4.7 Theoretical Results

I present a few theoretical results concerning the definitions presented in this chapter. Further theoretical results that involve work in both this and the next chapter are given at the end of that chapter.

I shall use the following symbols: Let $\mathcal{K}$ be a vocabulary, $\mathcal{G}_{\mathcal{K}}$ the set of literals, $\Gamma \subseteq \mathcal{G}_{\mathcal{K}}$ a set of facts, Θ a set of epistemic defeasible rules and Δ a DL ontology. Let V_P be the sets of positive and negative values for some patient p.

The $Rewrite$ function rewrites epistemic rules into ergonic. It seems sensible to enquire whether this process can be reversed, and I consider the general case first. A function is invertible if it is one-to-one and onto, but $Rewrite$ is not one-to-one; hence it is not invertible. Let $\mathcal{V}_p^+$ be some set of positive values from some patient p. Consider a pair of epistemic rules, r_1 and r_2 such that $Body(r_1) = Body(r_2)$ and $Head(r_1)$, $Head(r_2) \in \mathcal{V}_P^+$, but $Head(r_1) \neq Head(r_2)$. Then r_1 is distinct from r_2. However, $Rewrite(r_1) = Rewrite(r_2)$, and therefore are not distinct, and so $Rewrite$ is not one-to-one. I prove this with a counter-example

Proposition 4.7.1. *There does not exist an inverse of the function* $Rewrite$

Proof. Let r_a and r_b be two epistemic rules, such that:

r_a: People(x)
$\wedge$ hasTreatment(x,y) $\wedge$ TamoxifenTypes(y)
$\wedge$ IncreasedOS(z) $\wedge$ hasValue(z,1.2)
$\Rightarrow$ hasDeltaRisk(x,z)

r_b: People(x)
$\wedge$ hasTreatment(x,y) $\wedge$ TamoxifenTypes(y)
$\wedge$ IncreasedOS(z) $\wedge$ hasValue(z,1.4)
$\Rightarrow$ hasDeltaRisk(x,z)

and let $\mathcal{V}_p^+$ = {IncreasedOS(Increased1.4), IncreasedOS(Increased1.2)}.
Then r_a^* = $Rewrite(r_a)$ and r_b^* = $Rewrite(r_b)$, such that:
r_a^*: People(x)
$\wedge$ TamoxifenTypes(y)
$\Rightarrow$ hasPosIntent(x,y)

r_b^*: People(x)
$\wedge$ TamoxifenTypes(y)
$\Rightarrow$ hasPosIntent(x,y)

Then $Rewrite(r_a)$ = $Rewrite(r_b)$, and $r_a \neq r_b$, and hence $Rewrite$ is not invertible. □

The reason for the lack of invertibility is because $Rewrite()$ is 'lossy'. A change to the definition would allow it to be invertible, but I have chosen to keep the simpler option here. There is also a work-around, because each defeasible rule has a unique label, and so one could 'trace' a path back from the ergonic rule to the original epistemic rule, but this is not the same as the function being invertible.

One of the ideas behind the development of the the value-based rewriting function is that we can author a set of epistemic rules, based on the literature, and then use the $Rewrite()$ function to develop ergonic rules. Intuitively, there should be some relationship between a set of epistemic rules and the ergonic rules developed from them. I start by examining the most general case, where we see that not all epistemic rules can be rewritten to form ergonic ones.

Proposition 4.7.2. *There exists an epistemic rule r, such that for all sets of patient values, V_P^+ and V_p^-, $Rewrite(r) = null$.*

Proof. Note that the definition of $Rewrite()$ ensures that it only applies to epistemic rules whose head is of the form hasDeltaRisk(t_1, t_2) for $t_1, t_2 \in I$. Consider an epistemic rule whose head is of the form hasRisk(t_1, t_2). Then $Rewrite()$ does not apply to such a rule, and hence $Rewrite(r) = null$. □

However, if we restrict the contents of Θ to a particular form of epistemic rule only about changes in risk, rather than absolute risk, then we are able to prove we can always develop an ergonic rule from it. This is formalised below.

Proposition 4.7.3. *For every epistemic rule r where $Head(r)$ is of the form hasDeltaRisk(t_1, t_2), there exists a set of patient values, $\mathcal{V}_P^+$ and $\mathcal{V}_p^-$ such that $Head(Rewrite(r))$ = hasPosIntent(t_1, t_2) or $Head(Rewrite(r))$ = hasNegIntent(t_1, t_2)*

Proof. Since $\mathcal{V}_P^+$ or $\mathcal{V}_P^-$ are sets of literals of the form hasDeltaRisk(t_1, t_2), for any rule where $Head(r)$ is of the form hasDeltaRisk(t_1, t_2), there is some $\mathcal{V}_p^+$ or $\mathcal{V}_p^-$ such that $Head(r) \in \mathcal{V}_p^+$ or $Head(r) \in \mathcal{V}_p^-$. From the definition of $Rewrite()$, for every epistemic rule r where $Head(r) \in \mathcal{V}_P^+$ or $\mathcal{V}_P^-$, $Head(Rewrite(r))$ = hasPosIntent(t_1, t_2) or $Head(Rewrite(r))$ = hasNegIntent(t_1, t_2). □

Proposition 4.7.4. *For every set of epistemic rules $\Theta \subseteq \mathcal{R}_{\mathcal{K}}^{\mathcal{E}}$, and a corresponding set of ergonic rules, Φ = {$Rewrite(r|r \in \Theta$ and $Rewrite(r) \neq null)$ }, $|\Phi| \leq |\Theta|$.*

Proof. From inspection of $Rewrite$, we can see that for any epistemic rule $r \in \mathcal{R}_{\mathcal{K}}^{\mathcal{E}}$, the output of $Rewrite(r)$ is at most a single rule. Therefore for a set of epistemic rules $\Theta \subseteq \mathcal{R}_{\mathcal{K}}^{\mathcal{E}}$, there are at most as many ergonic rules produced in the output of $Rewrite(r) \mid r \in \Theta$ as there are epistemic rules in Θ. □

In general, therefore, we can be sure that for all epistemic rules where the head of the rule is of the form hasDeltaRisk(t_1, t_2), there is some set of patient values such that we can rewrite them to an ergonic rule. However, this rewriting function is not one-to-one, and so even for a set of such rules, the corresponding set of ergonic may be smaller than that of epistemic rules.

4.8 Summary

This chapter has covered the introduction and use of patient-specific values and intentions in constructing argument about courses of action. It has also shown that it captures some of our intuitions, and I have provided some brief theoretical results of the properties. However, as well as reasoning about what we might want to do, I also want to reason about the consequences of our actions, and this forms the basis of the next chapter.

Chapter 5

Reasoning about consequences

This chapter presents a technique for reasoning about the effects of carrying out an intention. I call this hypothetical reasoning, as it depends on a hypothesis that an intention to act will be fulfilled. As a consequence of hypothetical reasoning, we are able to use the information in epistemic defeasible rules to construct arguments about what might happen as the result of some treatment. The work in this chapter follows closely on from the last chapter, and extends the worked example given there. I conclude by presenting some theoretical results about the work in this chapter and the previous one, and then discuss this and the last chapter in the light of two previously published proposals. I briefly review these, and highlight some of the differences in approach between their work and mine. The novel work in this chapter lies in the definitions and theoretical results about the relationships between different types of argument.

5.1 From intention to hypothetical belief

So far I have been careful to observe an important distinction: Everything recorded in the A-Box of the ontology is thought to be true. The A-Box is a database to record what we know about the world: which patients there are, what their diseases are, etc. Following the work in the last chapter, it also contains a record of what we intend to do. In contrast, the work in this chapter allows us to construct arguments based on things that are not true when we construct the argument, but that may be true in the future. This needs some justification, and I shall start with the example from the beginning of the last chapter, where we had a rule of the form:

BMJ1999: People(x)
$\wedge$ hasDisease(x,y) $\wedge$ BreastCancerTypes(y)
$\wedge$ hasTreatment(x,z) $\wedge$ TamoxifenTypes(z)
$\wedge$ IncreasedOS(w)
$\Rightarrow$ hasDeltaRisk(x,w)

The point of the last chapter was to take these, along with information about valuing certain outcomes, to develop rules of the form:

BMJ1999*: People(x)

$\wedge$ hasDisease(x,y) $\wedge$ BreastCancerTypes(y)

$\wedge$ TamoxifenTypes(z)

$\Rightarrow$ hasPosIntent(x,z)

so that we can then develop an argument such as:

$\langle$BMJ1999*, hasPosIntent(MsJones,TamE5YrCourse) $\rangle$

to argue about whether or not we should give 5 years worth of tamoxifen.

However, as I noted in Secn. 6 of the last chapter, this allows us to argue about the intention to give a treatment, but not its effects, because the rules describing the effects of taking (for example) 5 years of tamoxifen have not been satisfied by the claim about intentions. To do that, we would need an argument of the form:

$\langle r_1$, hasTreatment(MsJones, TamE5YrCourse) $\rangle$

the claim of which would then satisfy the body of BMJ1999 in the example above. This would then let us use our existing knowledge to see what the effect of those actions would be, without having to re-write all the knowledge in our domain.

5.2 Committed Arguments

So far, ergonic arguments argue about an intention to perform some action. What I now want to do is to assume that an argument for an *intention* to do x is an argument *for believing that x has been done*. However, since these new arguments, which I shall refer to as *committed* arguments, are based upon the fact that we have an argument for the intention to perform an action in the first place, it would seem sensible to ensure that there remains some link between them. As we shall see, this is a matter of the presence of a sub-argument, and in this case is relatively easy to ensure.

I start by recalling that an ergonic argument is of the form $\langle A$, hasPosIntent(t_1, t_2)$\rangle$ (resp.$\langle A'$, hasNegIntent(t_1, t_2)$\rangle$). My aim is to be able to derive an argument of the form $\langle B$, hasTreatment(c,d)$\rangle$ (resp. $\langle B$, ¬hasTreatment(c,d)$\rangle$. Consider some argument a_1 = $\langle A$, hasPosIntent(t_1, t_2) $\rangle$. Then we want to make an argument of the form b_1 = $\langle B$, hasTreatment(t_1, t_2)$\rangle$. The problem with this approach is that there is no clear link between a_1 and b_1; therefore, if we later decide that a_1 is not warranted, that has no effect on b_1. This could lead us into a situation where we had many arguments about a treatment and its effects, without having any warranted arguments for the intention to perform the action itself. In order to avoid this problem, I develop a new type of defeasible rule, and add this to the support of a_1 in order to develop b_1.

Definition 5.2.1. ***Positive Commital Rules*** Let $\mathcal{K}$ be a vocabulary, I be the set of concrete individuals in $\mathcal{K}$, and $\mathcal{R}_{\mathcal{K}}$ be the set of defeasible rules. Let t_1, $t_2 \in I$ be concrete individuals such that $t_1 \in PeopleNames(\Omega)$ and $t_2 \in TreatmentTypes(\Omega)$. Then a defeasible rule $r \in \mathcal{R}_{\mathcal{K}}$ is a ***positive commital rule*** iff it is of the form hasPosIntent(t_1, t_2) $\Rightarrow$ hasTreatment(t_1,t_2). The set of all positive commital rules is denoted $\mathcal{R}_{\mathcal{K}}^{\mathcal{C}+}$.

Similarly, I define negative commital rules.

Definition 5.2.2. ***Negative Commital Rules*** Let $\mathcal{K}$ be a vocabulary, I be the set of concrete individuals in $\mathcal{K}$, and $\mathcal{R}_{\mathcal{K}}$ be the set of defeasible rules. Let t_1, $t_2 \in I$ be two concrete individuals such that $t_1 \in PeopleNames(\Omega)$ and $t_2 \in TreatmentTypes(t_2)$ hold. Then a defeasible rule $r \in \mathcal{R}_{\mathcal{K}}$ is a ***negative commital rule*** iff it is of the form hasNegIntent(t_1, t_2) $\Rightarrow$ ¬hasTreatment(t_1, t_2). The set of all negative commital rules is denoted $\mathcal{R}_{\mathcal{K}}^{\mathcal{C}\text{-}}$.

It will be useful to refer to both positive and negative commital rules.

Definition 5.2.3. A ***commital rule*** is either a positive commital rule or a negative commital rule. The set of all commital rules is denoted $\mathcal{R}_{\mathcal{K}}^{\mathcal{C}}$, such that $\mathcal{R}_{\mathcal{K}}^{\mathcal{C}} = \mathcal{R}_{\mathcal{K}}^{\mathcal{C}+} \cup \mathcal{R}_{\mathcal{K}}^{\mathcal{C}\text{-}}$.

Example 5.2.4. The following are both commital rules:

hasPosIntent(MsJones, TamE5YrCourse) $\Rightarrow$ hasTreatment(MsJones, TamE5YrCourse)

hasNegIntent(MsJones, TamD2YrCourse) $\Rightarrow$ ¬hasTreatment(MsJones, TamD2YrCourse)

My intention is to use commital rules to help develop new arguments by adding a commital rule to the support of an ergonic argument; this then creates an argument whose claim is that someone is or is not taking a treatment. Such arguments are termed committed arguments.

Definition 5.2.5. ***Committed Argument***

Let $\mathcal{K}$ be a vocabulary and $\mathcal{R}_{\mathcal{K}}$ be the set of defeasible rules. Let r_c be some commital rule in $\mathcal{R}_{\mathcal{K}}^{\mathcal{C}}$. Then an argument $\langle A_1, \phi_1 \rangle$ is a commited argument iff :

1. $\langle A_1, \phi_1 \rangle$ is an argument
2. There exists an ergonic argument $\langle A_2, \phi_2 \rangle$ that is a proper subargument of $\langle A_1, \phi_1 \rangle$
3. $Body(r_c) = Claim(\langle A_2, \phi_2 \rangle)$
4. $A_1 = A_2 \cup \{r_c\}$

The set of all committed arguments is denoted $\mathcal{A}_{\mathcal{K}}^{\mathcal{C}}$ such that $\mathcal{A}_{\mathcal{K}}^{\mathcal{C}} \subseteq \mathcal{A}_{\mathcal{K}}$

Now we have a definition for the committed arguments, I define a function to produce them:

Definition 5.2.6. Let $\mathcal{K}$ be some vocabulary, I be the set of concrete individuals, t_1, t_2 be elements of I and $\mathcal{R}_{\mathcal{K}}$ the set of defeasible rules. Let $e \in \mathcal{A}^{\mathcal{R}}$ be some ergonic argument, $r_c^+ \in \mathcal{R}_{\mathcal{K}}^{\mathcal{C}+}$ be a positive commital rule and $r_c^- \in \mathcal{R}_{\mathcal{K}}^{\mathcal{C}-}$ a negative commital rule, such that $Body(r_c^+) = Claim(e)$ and $Body(r_c^-) = Claim(e)$. The $Commit : \mathcal{A}^{\mathcal{R}}_{.} \mapsto \mathcal{A}^{\mathcal{C}}_{.}$ function takes an ergonic argument and returns a commited one, so that for some ergonic argument e:

If e is of the form $\langle \{r_1...r_n\}$, hasPosIntent(t_1,t_2) $\rangle$, then $Commit(e)$, is of the form

$\langle \{r_1...r_n\} \cup \{r_c^+\}$, hasTreatment($t_1$,$t_2$)$\rangle$

If e is of the form $\langle \{r_1...r_n\}$, hasNegIntent(t_1,t_2) $\rangle$, then $Commit(e)$, is of the form

$\langle \{r_1...r_n\} \cup \{r_c^-\}$, ¬hasTreatment($t_1$,$t_2$)$\rangle$

The advantage of this approach is that e is clearly a subargument of $Commit(e)$. Thus any arguments that attacks e will attack $Commit(e)$ (see 5.4.14 for a proof of this).

5.3 Hypothetical Arguments

So far we have seen how we can develop commited arguments from ergonic ones. However, the motivation for this was to be able to use our existing knowledge about the effects of treatment in order to reason about the effects of giving possible treatments. In order to do this, we can use the existing epistemic rules. By doing so, we then allow ourselves to reason about the possible outcomes of the action, using arguments that are epistemic in nature. Not only does this satisfy my criteria from chapter 1 (1.4), in that we can now reason about an action and its outcomes, but since we move back to epistemic reasoning, it allows for the potential to develop another round of ergonic-committed-epistemic reasoning. However, the arguments we develop are different to the purely epistemic arguments that we develop based on facts about treatment, as they only exist as a result of the process of commital. Because of this it will be useful to distinguish those epistemic arguments whose support comes entirely from $\mathcal{R}_{\mathcal{K}}^{\mathcal{P}}$ and those developed as a result of ergonic reasoning and commital. The first group are purely epistemic arguments; I shall call the latter sort hypothetical arguments - they are those whose support contains epistemic rules, but which have ergonic and committed subarguments.

Definition 5.3.1. Let $\mathcal{K}$ be some vocabulary, $\mathcal{G}_{\mathcal{K}}$ the set of literals and I the set of concrete individuals in $\mathcal{K}$. Let $t_1, t_2 \in I$ be concrete individuals and let $\Gamma \subseteq \mathcal{G}_{\mathcal{K}}$ be a set of literals such that $\Delta \cup \Gamma \vdash_{Ont}$ People(t_1) and $\Delta \vdash_{Ont}$ TreatmentTypes(t_2). Let $\mathcal{R}_{\mathcal{K}}^{\mathcal{P}}$ be the set of all epistemic rules, $\mathcal{R}_{\mathcal{K}}^{\mathcal{R}}$ the set of all ergonic rules, $\mathcal{R}_{\mathcal{K}}^{\mathcal{C}}$ the set of all commital rules, and let $r_p \in \mathcal{R}_{\mathcal{K}}^{\mathcal{P}}$ be an epistemic rule. Let $\Omega = (\Theta, \Delta)$ be an OAF. A ***hypothetical argument*** is an argument$\langle A_1, \phi_1\rangle$ from Ω where:

1. $\langle A_1, \phi_1\rangle$ is an argument
2. There is a committed argument $\langle A_2, \phi_2 \rangle$ which is a proper subargument of $\langle A_1, \phi_1\rangle$
3. $A_1 = A_2 \cup \{r_p\}$

The set of all hypothetical arguments is denoted $\mathcal{A}_{\mathcal{K}}^{\mathcal{H}}$.

Example 5.3.2. I will reuse the worked example from the end of Chapter four, which I recall here:

Let $\mathcal{K}$ be a vocabulary, $\mathcal{G}_{\mathcal{K}}$ the set of literals, $\Gamma \subseteq \mathcal{G}_{\mathcal{K}}$ a set of facts, Θ a set of epistemic defeasible rules and Δ_{BC} the breast cancer ontology. Let $V_{MsJones}$ be the sets of positive and negative values for MsJones, and Φ be set of ergonic defeasible rules. Then $\Omega = (\Theta \cup \Phi$, $\Delta_{BC} \cup \Gamma$, $V_p)$ is an E-OAF, where Θ, Φ, Γ and $V_{MsJones}$ are as below.

Γ = {Woman(MsJones), hasDisease(MsJones,ProtoBreastCancer)}

$\mathcal{V}_{MsJones}^{+}$ ={hasDeltaRisk(MsJones,IncreasedOS1.01),

hasDeltaRisk(MsJones,IncreasedOS1.02),

hasDeltaRisk(MsJones,IncreasedOS1.03),

hasDeltaRisk(MsJones,IncreasedOS1.1),

hasDeltaRisk(MsJones,IncreasedOS1.2)}

and

$\mathcal{V}^{-}_{MsJones}$ ={hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer2.3),

hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer2.53),

hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer4.1),

hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer6.0),

hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer6.4),

hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer7.5)

hasDeltaRisk(MsJones, DecreasedOS0.7)}.

$V_{MsJones} = (\mathcal{V}^{+}_{MsJones}, \mathcal{V}^{-}_{MsJones})$

Θ = {NICE2003, NLCN1997, NCRN2001, BMJ1999, BJC2004}

where

NICE2003: People(x)

$\wedge$ hasDisease(x,y) $\wedge$ BreastCancerTypes(y)

$\wedge$ Tamoxifen5YrCourse(z)

$\Rightarrow$ hasTreatment(x,z)

NLCN1997: People(x)

$\wedge$ hasDisease(x,y) $\wedge$ BreastCancerTypes(y)

$\wedge$ Tamoxifen2YrCourse(z)

$\Rightarrow$ hasTreatment(x,z)

NCRN2001: People(x)

$\wedge$ hasDisease(x,y) $\wedge$ BreastCancerTypes(y)

$\wedge$ ECMFChemotherapy(z)

$\Rightarrow$ hasTreatment(x,z)

BMJ1999: People(x)

$\wedge$ hasTreatment(x,y) $\wedge$ TamoxifenByLength(y)

$\wedge$ IncreasedOS(z) $\wedge$ hasValue(z,1.2)

$\Rightarrow$ hasDeltaRisk(x,z)

BJC2004: People(x)

$\wedge$ hasTreatment(x,y) $\wedge$ TamoxifenByLength(y)

$\wedge$ IncreasedRiskEndometrialCancer(z) $\wedge$ hasValue(z,4.1)

$\Rightarrow$ hasDeltaRisk(x,z)

and

$\Phi = \{Rewrite(r) \mid r \in \Theta$ and $Rewrite(r) \neq null\}$

and we have the following commital rules:

r^+_{c1} : hasPosIntent(MsJones, TamD2YrCourse) $\Rightarrow$ hasTreatment(MsJones, TamD2YrCourse)

r^+_{c2} : hasPosIntent(MsJones, TamE5YrCourse) $\Rightarrow$ hasTreatment(MsJones, TamE5YrCourse)

r^+_{c3} : hasPosIntent(MsJones, TamFgt5YrCourse) $\Rightarrow$ hasTreatment(MsJones, TamFgt5YrCourse)

r^-_{c1} : hasNegIntent(MsJones, TamD2YrCourse) $\Rightarrow$ $\neg$hasTreatment(MsJones, TamD2YrCourse)

r^-_{c2} : hasNegIntent(MsJones, TamE5YrCourse) $\Rightarrow$ $\neg$hasTreatment(MsJones, TamE5YrCourse)

r^-_{c3} : hasNegIntent(MsJones, TamFgt5YrCourse) $\Rightarrow$ $\neg$hasTreatment(MsJones, TamFgt5YrCourse)

We form the following arguments from Ω:

a_1 $\langle$\{NICE2003\}, hasTreatment(MsJones,TamE5YrCourse)$\rangle$

a_2 $\langle$\{NLCN1997\}, hasTreatment(MsJones,TamD2YrCourse)$\rangle$

a_3 $\langle$\{NCRN2001\}, hasTreatment(MsJones,ECMFTypeInstance)$\rangle$

a_4 $\langle$\{BMJ1999$^*_1$\}, hasPosIntent(MsJones,TamD2YrCourse)$\rangle$

a_5 $\langle$\{BMJ1999$^*_2$\}, hasPosIntent(MsJones,TamE5YrCourse)$\rangle$

a_6 $\langle$\{BMJ1999$^*_3$\}, hasPosIntent(MsJones,TamFgt5YrCourse)$\rangle$

a_7 $\langle$\{BJC2004$^*_1$\}, hasNegIntent(MsJones,TamD2YrCourse)$\rangle$

a_8 $\langle$\{BJC2004$^*_2$\}, hasNegIntent(MsJones,TamE5YrCourse)$\rangle$

a_9 $\langle$\{BJC2004$^*_3$\}, hasNegIntent(MsJones,TamFgt5YrCourse)$\rangle$

a_{10} $\langle$\{NICE2003*, BMJ1999\}, hasDeltaRisk(MsJones, IncreasedOS1.1)$\rangle$

a_{11} $\langle$\{NLCN1997*, BMJ1999\}, hasDeltaRisk(MsJones, IncreasedOS1.1)$\rangle$

a_{12} $\langle$\{NICE2003*, BJC2004\}, hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer2.3)$\rangle$

a_{13} $\langle$\{NLCN1997*, BJC2004\}, hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer2.3)$\rangle$

and then can form the following committed arguments:

a_{14} $\langle$\{BMJ1999$^*_1$, r^+_{c1}\}, hasTreatment(MsJones,TamD2YrCourse)$\rangle$

a_{15} $\langle$\{BMJ1999$^*_2$, r^+_{c2}\}, hasTreatment(MsJones,TamE5YrCourse)$\rangle$

a_{16} $\langle$\{BMJ1999$^*_3$, r^+_{c3}\}, hasTreatment(MsJones,TamFgt5YrCourse)$\rangle$

a_{17} $\langle$\{BJC2004$^*_1$, r^-_{c1}\}, $\neg$hasTreatment(MsJones,TamD2YrCourse)$\rangle$

a_{18} $\langle$\{BJC2004$^*_2$, r^-_{c2}\}, $\neg$hasTreatment(MsJones,TamE5YrCourse)$\rangle$

a_{19} $\langle$\{BJC2004$^*_3$, r^-_{c3}\}, $\neg$hasTreatment(MsJones,TamFgt5YrCourse)$\rangle$

and hypothetical arguments:

a_{20} $\langle${BMJ1999$^*_1$, r^+_{c1}, BMJ1999}, hasDeltaRisk(MsJones, IncreasedOS1.1)$\rangle$

a_{21} $\langle${BMJ1999$^*_2$, r^+_{c2}, BMJ1999}, hasDeltaRisk(MsJones, IncreasedOS1.1)$\rangle$

a_{22} $\langle${BMJ1999$^*_3$, r^+_{c3}, BMJ1999}, hasDeltaRisk(MsJones, IncreasedOS1.1)$\rangle$

a_{23} $\langle${BMJ1999$^*_1$, r^+_{c1}, BMJ1999}, hasDeltaRisk(MsJones, IncreasedRiskEndometrial-Cancer2.3)$\rangle$

a_{24} $\langle${BMJ1999$^*_2$, r^+_{c2}, BMJ1999}, hasDeltaRisk(MsJones, IncreasedRiskEndometrial-Cancer2.3)$\rangle$

a_{25} $\langle${BMJ1999$^*_3$, r^+_{c3}, BMJ1999}, hasDeltaRisk(MsJones, IncreasedRiskEndometrial-Cancer2.3)$\rangle$

5.3.1 Informal analysis

At this point, we should note that a_{14} and a_{15} above are familiar. If we look at some of the arguments that we generated at the end of the last chapter (i.e. before I introduced committed and hypothetical reasoning):

a_{10} $\langle${NICE2003, BMJ1999}, hasDeltaRisk(MsJones, IncreasedOS1.2)$\rangle$

a_{11} $\langle${NLCN1997, BMJ1999}, hasDeltaRisk(MsJones, IncreasedOS1.2)$\rangle$

a_{12} $\langle${NICE2003, BJC2004}, hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer4.1)$\rangle$

a_{13} $\langle${NLCN1997, BJC2004}, hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer4.1)$\rangle$

they have the same claims as the arguments at the end of the example of above:

a_{20} $\langle${BMJ1999$^*_1$, r^+_{c1}, BMJ1999}, hasDeltaRisk(MsJones, IncreasedOS1.1)$\rangle$

a_{23} $\langle${BMJ1999$^*_1$, r^+_{c1}, BMJ1999}, hasDeltaRisk(MsJones, IncreasedRiskEndometrial-Cancer2.3)$\rangle$

If we examine the support of the arguments formed from our hypothetical reasoning (a_{20} - a_{25}), we see that it is {BMJ1999*, r^+_{c1-3}, BMJ1999}. In contrast, the arguments from earlier have supports with either NICE2003 or NLCN1997. Examining the rules is instructive.

NICE2003: People(x)
$\wedge$ hasDisease(x,y) $\wedge$ BreastCancerTypes(y)
$\wedge$ Tamoxifen5YrCourseTypes(z)
$\Rightarrow$ hasTreatment(x,z)

NLCN1997: People(x)
$\wedge$ hasDisease(x,y) $\wedge$ BreastCancerTypes(y)
$\wedge$ Tamoxifen2YrCourseTypes(z)
$\Rightarrow$ hasTreatment(x,z)

These are heuristic rules. Although they may seem sensible, such rules are quite different to those developed from clinical trials. Such rules capture one aspect of medical knowledge, and were useful as early examples of what I was trying to do, but if we look at what they say, they say that 'People with breast cancer take tamoxifen' - such rules might come from a treatment guideline. In contrast, BMJ1999 and BJC2004 are very different:

BMJ1999: People(x)
∧ hasTreatment(x,y) ∧ Tamoxifen(y)
∧ IncreasedOS(z)
⇒ hasOutcome(x,z)

BJC2004: People(x)
∧ hasTreatment(x,y) ∧ Tamoxifen(y)
∧ EndometrialCaTypes(z)
⇒ hasDisease(x,z)

These second two rules (BMJ1999 and BJC2004) (allowing for some simplifications) are much closer to what we would expect as the result of clinical trials. The similarity of the two types of argument is reassuring - what I have done is demonstrated how OAF can take rules from clinical trials, combine them with patient values and make the same claims as a guideline, but without the implicit assumptions that are contained in a guideline. This is not a proof that the two approaches are equivalent, but it is encouraging that we seem to be going in the right direction.

5.4 The effects of Hypothetical Reasoning

In the course of the last two chapters, I have presented a technique for using patient values to help generate arguments about courses of action, and have then shown how to derive the results of these actions. I have also presented a worked example, running over both chapters, to illustrate the effect of these ideas. I shall now add some further definitions, and then develop proofs about properties of argument composition and interaction. In order to do this, I provide a simple typology of arguments which matches that which we have already seen for rules (epistemic, ergonic, etc.). Recall that the definitions of hypothetical and committed arguments were given above.

5.4.1 Set relations

Definition 5.4.1. A set of rules, R, is ***epistemic*** iff it contains an epistemic rule***, ergonic*** if it contains an ergonic rule and ***heuristic*** if it contains a heuristic rule. R is said to be ***purely epistemic*** if every element of R is epistemic, ***purely ergonic*** if every element of R is ergonic, and ***purely heuristic*** if every element of R is heuristic.

We then use these definitions of sets of rules, together with the definitions earlier in the chapter, to complete the definition of types of argument. I do this by proving and disproving various relationships between different types of argument. Although somewhat mechanical, these results are useful in that they allow us to be certain about how the different types of argument relate to one another.

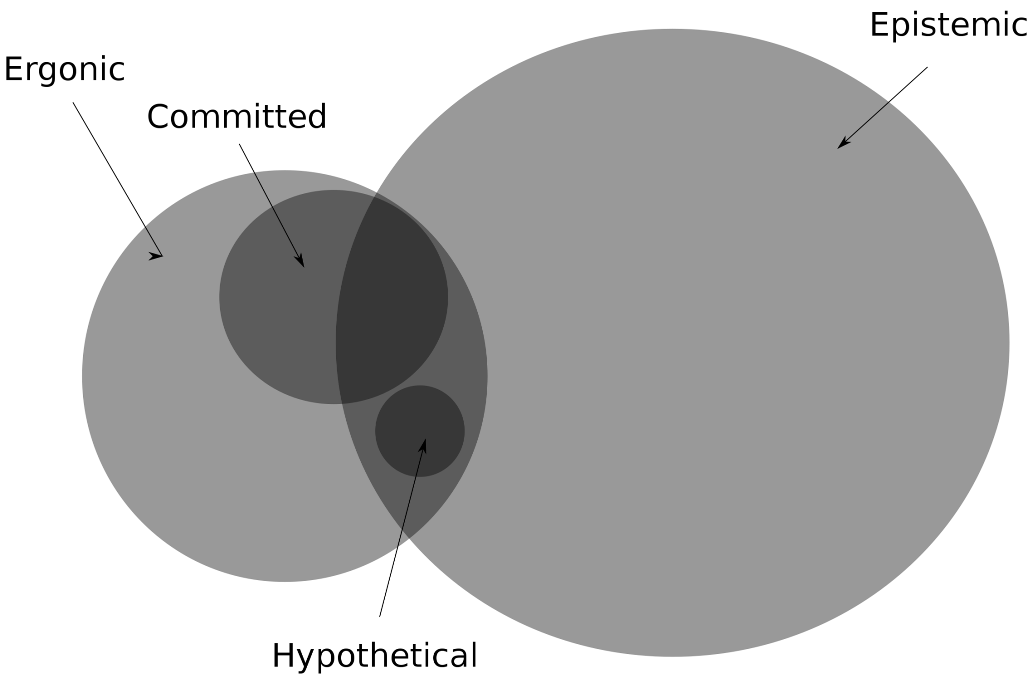

Figure 5.1: Venn diagram of Argument Types. The relationships follow the proofs given in Secn. 5.4.1

Definition 5.4.2. An argument A is epistemic (purely epistemic) if $Support(A)$ is epistemic (purely epistemic), is ergonic (purely ergonic) if $Support(A)$ is ergonic (purely ergonic) and heuristic (purely heuristic) if $Support(A)$ is heuristic (purely heuristic). The set of epistemic arguments is denoted $\mathcal{A}^{\mathcal{P}}_{\mathcal{K}}$, that of ergonic arguments is denoted $\mathcal{A}^{\mathcal{R}}_{\mathcal{K}}$ and that of heuristic arguments $\mathcal{A}^{\mathcal{U}}_{\mathcal{K}}$. The sets of commited and hypothetical arguments are denoted $\mathcal{A}^{\mathcal{C}}_{\mathcal{K}}$ and $\mathcal{A}^{\mathcal{H}}_{\mathcal{K}}$ as above.

It would be tempting to suggest that there is no argument that is both epistemic and ergonic - that is, $\mathcal{A}^{\mathcal{P}}_{\mathcal{K}} \cap \mathcal{A}^{\mathcal{R}}_{\mathcal{K}} = \emptyset$. This is not the case.

Example 5.4.3. In the worked example above, a_4- a_9 are purely ergonic and a_{14} - a_{19} are committed, while a_{20}- $_{25}$ are jointly epistemic, ergonic and hypothetical.

In order to understand the relationship between the various argument types, I provide a diagram showing the relationship between the sets. This diagram is developed in a stepwise fashion below. I start by summarising the notation I shall use. Let $\mathcal{K}$ be some vocabulary, I the set of concrete individuals in $\mathcal{K}$, $t_1, t_2 \in I$ be concrete individuals, Δ some DL ontology, $\mathcal{R}^{\mathcal{P}}_{\mathcal{K}}$ the set of epistemic rules, $\mathcal{R}^{\mathcal{R}}_{\mathcal{K}}$ the set of ergonic rules and $\mathcal{R}^{\mathcal{C}}_{\mathcal{K}}$ the set of commital rules, with r_c^+ a positive commital rule and r_c^- a negative commital rule. Let $\mathcal{A}_{\mathcal{K}}$ be the set of all arguments, while $\mathcal{A}^{\mathcal{P}}_{\mathcal{K}}$, $\mathcal{A}^{\mathcal{R}}_{\mathcal{K}}$, $\mathcal{A}^{\mathcal{U}}_{\mathcal{K}}$, $\mathcal{A}^{\mathcal{C}}_{\mathcal{K}}$ and $\mathcal{A}^{\mathcal{H}}_{\mathcal{K}}$ are as defined above.

I shall proceed by proving that various relationships hold between these sets, and then use this information to develop a diagram. I have already shown that $\mathcal{A}^{\mathcal{P}}_{.}$ and $\mathcal{A}^{\mathcal{R}}_{.}$ are not disjoint; I now show that neither is a subset of the other. In order to do so, I show that the sets of pure arguments are disjoint.

Proposition 5.4.4. *The sets of pure ergonic and pure epistemic arguments are disjoint*

Proof. Let I be the set of named individuals in Δ and $t_1, t_2 \in I$. Let r be a defeasible rule, and R a set of defeasible rules. Recall that an epistemic rule is defined such that:

1. $Head(r)$ = hasDeltaRisk(t_1,t_2) or $Head(r)$ = hasRisk(t_1,t_2)} and
2. There is at most one atom of the form hasTreatment(t_1,t_3) $\in Body(r)$

and ergonic rules are such that:

1. $Head(r)$ = hasPosIntent(t_1,t_2) or
2. $Head(r)$ = hasNegIntent(t_1,t_2)}

No rule is both ergonic and epistemic. Therefore by def. 5.4.1, if R is purely epistemic, then R is not ergonic and if R is purely ergonic, then R is not epistemic. So for an argument a :$\langle A, \phi\rangle$, if a is purely epistemic then a is not ergonic, and if a is purely ergonic it is not epistemic. Therefore the sets of such arguments are disjoint. □

As we can see, while an arguments may be both ergonic and epistemic, there are arguments that are only one or only the other. Thus $\mathcal{A}^{\mathcal{P}}_{\cdot} \cap \mathcal{A}^{\mathcal{R}}_{\cdot} \neq \emptyset$, but $\mathcal{A}^{\mathcal{P}}_{\cdot} \not\subseteq \mathcal{A}^{\mathcal{R}}_{\cdot}$ and $\mathcal{A}^{\mathcal{R}}_{\cdot} \not\subseteq \mathcal{A}^{\mathcal{P}}_{\cdot}$. However, in the definition of a commited argument above, I suggested that this was not the case with commited and ergonic arguments. I now prove this.

Proposition 5.4.5. *Every committed argument is also ergonic*

Proof. This is trivially true from the definition of a committed argument: every committed argument has an ergonic subargument, and if $Support(e)$ is ergonic, $Support(e) \cup r$ is ergonic. □

Proposition 5.4.6. *Every hypothetical argument is also ergonic*

Proof. This is trivially true from the definition of a hypothetical argument: every hypothetical argument has a committed subargument (which is ergonic, from the proof above) and if $Support(c)$ is ergonic, $Support(c) \cup r$ is ergonic. □

Proposition 5.4.7. *There is no argument that is both hypothetical and committed*

Proof. If r is an epistemic rule, then $Head(r)$ = hasRisk(t_1, t_2) or $Head(r)$ = hasDeltaRisk(t_1, t_2). If r is a commital rule then $Head(r)$ = hasTreatment(t_1, t_2) or ¬hasTreatment(t_1, t_2). Therefore there is no rule that is both epistemic and a commital rule, and hence the sets of such rules are disjoint. For an argument c_1= $\langle A_1, \phi_1\rangle$ to be a commited argument, there must be an argument e = $\langle A_2, \phi_2\rangle$ that is a proper subargument of c, and where there is a commital rule $r_c \in \mathcal{R}^{\mathcal{C}}_{\mathcal{K}}$ such that A_1 = $A_2 \cup r_c$. For an argument h = $\langle A_3, \phi_3\rangle$ to be a hypothetical argument, there must be a commited argument c_2 = $\langle A_4, \phi_4\rangle$ which is a proper subargument of h, and where there is an epistemic rule $r_r \in \mathcal{R}^{\mathcal{R}}_{\mathcal{K}}$ such that A_3 = $A_4 \cup$

r_r. Since all committed arguments are ergonic, it could be the case that the ergonic subargument of c_1 is also a commited argument, that is $e = c_2$. But for $h = c_1$, it would have to be the case that $r_r = r_c$, and as we have seen above, this is not possible. □

Proposition 5.4.8. *All hypothetical arguments have ergonic and committed arguments as proper sub-arguments*

Proof. Let $h \in \mathcal{A}^{\mathcal{H}}_{\mathcal{K}}$ be a hypothetical argument. Then from 5.4.6 and 5.4.7 above, there exists a committed argument $c \in \mathcal{A}^{\mathcal{C}}_{\mathcal{K}}$ and an ergonic argument $e \in \mathcal{A}^{\mathcal{R}}_{\mathcal{K}}$ such that $Support(e) \subset Support(c)$ and $Support(c) \subset Support(h)$. Therefore $Support(e) \subset Support(h)$ and thus e is a sub-argument of h. □

Proposition 5.4.9. *The sets of committed and hypothetical arguments are proper subsets of the set of ergonic arguments.*

Proof. Note that Proposition 5.4.8 implies that all committed arguments ($\mathcal{A}^{\mathcal{C}}_{\mathcal{K}}$) and hypothetical arguments ($\mathcal{A}^{\mathcal{H}}_{\mathcal{K}}$) arguments are ergonic. Therefore we have $\mathcal{A}^{\mathcal{C}}_{\mathcal{K}} \subseteq \mathcal{A}^{\mathcal{R}}_{\mathcal{K}}$ and $\mathcal{A}^{\mathcal{H}}_{\mathcal{K}} \subseteq \mathcal{A}^{\mathcal{R}}_{\mathcal{K}}$. The worked examples contain many ergonic arguments that are neither committed nor hypothetical, hence $\mathcal{A}^{\mathcal{R}}_{\mathcal{K}} \not\subset \mathcal{A}^{\mathcal{C}}_{\mathcal{K}}$ and $\mathcal{A}^{\mathcal{R}}_{\mathcal{K}} \not\subset \mathcal{A}^{\mathcal{H}}_{\mathcal{K}}$. Therefore $\mathcal{A}^{\mathcal{C}}_{\mathcal{K}} \subset \mathcal{A}^{\mathcal{R}}_{\mathcal{K}}$ and $\mathcal{A}^{\mathcal{H}}_{\mathcal{K}} \subset \mathcal{A}^{\mathcal{R}}_{\mathcal{K}}$ both hold. □

So far I have shown that the sets of epistemic and ergonic arguments intersect, and that all committed and hypothetical arguments are ergonic. I now examine the relationship between epistemic, committed and hypothetical arguments.

Proposition 5.4.10. *Every hypothetical argument is epistemic*

Proof. Consider the definition of a hypothetical argument. Then for some argument $h = \langle A, \phi\rangle$, such that h is hypothetical, then $Support(h)$ contains an epistemic rule, and hence h is epistemic. □

The next consideration is the relationship of committed and epistemic arguments. However, describing this is more difficult than that for hypothetical arguments, as I show below.

Example 5.4.11. Consider some commited argument $c_1 = \langle\{r^*_1, r^+_c\}, \phi\rangle$, where $r^*_1 \in \mathcal{R}^{\mathcal{R}}_{\mathcal{K}}$ is an ergonic rule, and $r^+_c \in \mathcal{R}^{\mathcal{C}}_{\mathcal{K}}$ is a commital rule. Then c_1 is not epistemic. However, if we consider some other argument $c_2 = \langle\{r^*_1, r^+_c, r_1\}, \phi\rangle$, where r^*_1 and r_c are as before and $r_1 \in \mathcal{R}^{\mathcal{P}}_{\mathcal{K}}$, then c_2 is epistemic.

Clearly therefore some committed arguments are epistemic, and some are not. However, it is useful to see if we can refine this further. From the example above, we might assume that every committed argument that is also epistemic is hypothetical. This is not true, as the following example shows:

Example 5.4.12. Consider a hypothetical argument $a_1 = \langle\{r^*_1, r^+_c, r_1\}, \phi_1\rangle$, where $r^*_1 \in \mathcal{R}^{\mathcal{R}}_{\mathcal{K}}$ is an ergonic rule, $r^+_c \in \mathcal{R}^{\mathcal{C}}_{\mathcal{K}}$ and $r_1 \in \mathcal{R}^{\mathcal{P}}_{\mathcal{K}}$, and $a_2 = \langle\{r^*_1, r^+_c\}, \phi_2\rangle$ is a commited subargument of a_1. Then a_1 is committed, epistemic and hypothetical. However, consider an argument $a_3 = \langle\{r^*_2, r^+_c, r_2\}, \phi\rangle$, which has a subargument $a_4 = \langle\{r^*_2, r_2\}, \phi\rangle$ which is ergonic and epistemic. Then a_3 is ergonic, epistemic and committed, but not hypothetical.

This is a suitable point at which to comment on the difference between committed and hypothetical arguments as compared to epistemic and ergonic arguments. Whereas epistemic and ergonic arguments are defined in terms the set of rules in the support, committed and hypothetical arguments are defined in terms of the *last* rule used in the last step of reasoning. I would suggest that they seem to capture some informal intuitions about when we would use such types of reasoning.

5.4.2 Attack relations

So far, we have considered argument relationships purely in terms of the type of arguments. However, the attack relationships between arguments are clearly important, and I now consider whether we can prove that such a relationship does not hold between certain types of arguments. To do so, I restrict myself to considering only foreground arguments, as given a set of rules, we can be certain about the claims of such argument, as I now show. Recall that the head of an ergonic rule is always of the form hasPosIntent(t_1, t_2) or hasNegIntent(t_1, t_2). Then we might assume that the claim of a pure ergonic argument is always of such a form, but as this example shows, this is not the case.

Example 5.4.13. Let $\mathcal{K}$ be a vocabulary, $\mathcal{G}_{\mathcal{K}}$ the set of literals and Δ a DL ontology such that there is an axiom in Δ^T of the form:

PeoplePosTamIntent $\equiv$ People $\sqcap$ $\exists$hasPosIntent. Tamoxifen

Let Γ be a set of facts, and Ω = (Θ, Δ) a S-OAF where Γ and Θ are as below:

Γ = {Women(MsJones), hasDisease(ProtoBreastCancer)}

Θ = {r_1}

where :

r_1: People(x)
$\wedge$ hasDisease(x,y) $\wedge$ BreastCancer(y)
$\wedge$ Tamoxifen(z)
$\Rightarrow$ hasPosIntent(y,z)

Then there is an argument of the form $\langle$ {r_1}, hasPosIntent(MsJones, TamE5YrCourse)$\rangle$, as well as one of the form $\langle$ {r_1}, PeoplePosTamIntent(MsJones)$\rangle$. r_1 is ergonic, and so both arguments are pure ergonic arguments, but the second is not a foreground argument. The problem is that in general, the T-Box of the ontology can contain arbitrary consistent formulae, and so we cannot be sure what the claims of background arguments may be. For that reason, we consider only foreground arguments.

Proposition 5.4.14. *Let e be an ergonic argument. Then any argument that attacks e also attacks $Commit(e)$*

Proof. Let e be an ergonic argument and c be the commited argument formed from e such that c = $Commit(e)$. Then $Support(c) = Support(e) \cup \{r_c\}$, and hence $Support(e) \subset Support(c)$. Then e is a subargument of c . Therefore, any argument a that attacks e also attacks c at e. $\square$

Proposition 5.4.15. *Let $t_1, t_2 \in I$. If e_1= $\langle A_1$,$\phi\rangle$ is a pure foreground ergonic argument, then ϕ =* hasPosIntent(t_1, t_2) *or ϕ =* hasNegIntent(t_1, t_2)

Proof. By definition, a pure ergonic argument is one whose support consists entirely of ergonic rules, and a foreground argument is one whose claim is the head of a rule. From the definition of an ergonic rule, all ergonic rules have a head of the form hasPosIntent(t_1,t_2) or hasNegIntent(t_1,t_2). Therefore, if the support of the argument e_1 consists entirely of such rules, then from the definition of a foreground argument, the claim e_1 must be of the form hasPosIntent(t_1, t_2) or hasNegIntent(t_1, t_2).

□

Proposition 5.4.16. *A pure foreground ergonic argument cannot be used as a counter-argument to a pure foreground epistemic argument*

Proof. I start by recalling the definition of a counter-argument and disagreement :

> We say that $\langle A_1, \phi_1\rangle$ counter-argues, rebuts or attacks $\langle A_2, \phi_2\rangle$ at ϕ iff there exists a sub-argument $\langle A, \phi\rangle$ of $\langle A_2, \phi_2\rangle$ such that ϕ and ϕ_1 disagree.
>
> We say that two ground formulae ϕ and $\phi_1 \in \mathcal{F}_\mathcal{K}$ ***disagree*** iff $conflict_\Delta(\phi, \phi_1)$ is true.

Let $t_1, t_2 \in I$. Consider some pure epistemic argument $\langle A_2, \phi_2\rangle$; then by definition, $Support(\langle A_2, \phi_2\rangle)$ contains only epistemic rules. From the definition of an epistemic rule, no such rule contains a literal of the form hasPosIntent(t_1, t_2) or hasNegIntent(t_1, t_2) (in either body or head). Since we are interested in pure epistemic arguments, there is no sub-argument $\langle A, \phi\rangle \subseteq \langle A_2, \phi_2\rangle$ whose claim is of such a form. If $\langle A_1, \phi_1\rangle$ is a pure ergonic argument, then ϕ_1 is either of the form hasPosIntent(t_1, t_2) or hasNegIntent(t_1, t_2) (by the proof above). Let ψ_1 = hasPosIntent(t_1, t_2) and ψ_2 = hasNegIntent(t_1, t_2). Then for some formulae $\phi \in \mathcal{F}_\mathcal{K}$, $conflict_\Delta(\psi_1$,$\phi_1)$ only holds if ϕ_1 = hasNegIntent(t_1, t_2) and $conflict_\Delta(\psi_2$,$\phi_2)$ only holds if ϕ_2 = hasPosIntent(t_1, t_2). Since the head of a purely ergonic argument is of the form given by either ψ_1 or ψ_2, and neither ϕ_1 nor ϕ_2 appear in an epistemic argument, the two argument types cannot conflict. □

Because arguments can be of multiple types this does not hold for arguments that are not pure, as we see here.

Example 5.4.17. Let r_1- r_3 be defeasible rules such that $r_3 \in \mathcal{R}^\mathcal{P}_\mathcal{K}$, $r_1, r_2 \in \mathcal{R}^\mathcal{R}_\mathcal{K}$ and let $r_c \in \mathcal{R}^\mathcal{C}_\mathcal{K}$ be a commital rule. Let $\phi_1 ... \phi_3$ be literals in $\mathcal{F}_\mathcal{K}$ and Δ be some DL ontology, such that $conflict_\Delta(\phi_1, \phi_2)$ holds. Consider an argument a_1 = $\langle\{r_1, r_c, r_3\}$,$\phi_3\rangle$ and another a_2 = $\langle\{r_2\}$,$\phi_2\rangle$. Then a_1 is both an epistemic argument and an ergonic one, and a_2 is a pure ergonic argument. Since $r_1 \in \mathcal{R}^\mathcal{P}_\mathcal{K}$, there is a ergonic subargument $\langle\{r_1\}$,$\phi_1\rangle \subset \langle\{r_1, r_c, r_3\}$,$\phi_3\rangle$. Since $conflict_\Delta(\phi_1, \phi_2)$ holds, a_2 counter-argues a_1 at ϕ_1.

In general, because arguments may be of multiple types, it is difficult to make general assertions about types of arguments in general. However, we can make an exception for pure ergonic arguments where the rules used in the support come solely from rewriting epistemic rules. It will be useful to define such rules.

Definition 5.4.18. Let $\Omega = (\Theta, \Delta, V_p)$ be an E-OAF. An ergonic defeasible rule, r_r, is a ***derived ergonic*** defeasible rule iff r_r is an ergonic rule and there exists some epistemic rule $r_p \in \Theta$ such that $r_r = Rewrite(r_p)$.

The definition of a derived ergonic rule captures the situation where our knowledge-base is largely, or entirely epistemic in nature, and ergonic rules are derived by rewriting epistemic rules based on a patient's values. This is true for most of the work in the thesis.

Proposition 5.4.19. *Let $e = \langle A, \phi \rangle$ be a pure ergonic argument from Ω, such that every element of A is also a derived ergonic rule. Then $Size(a)$ = 1*

Proof. Recall that for an epistemic rule $r_i \in \mathcal{R}_{\mathcal{K}}^{\mathcal{P}}$, hasPosIntent($t_1$,$t_2$) $\notin Atoms(r)$ and hasNegIntent(t_1,t_2) $\notin Atoms(r)$. Then for all derived ergonic rules $r_i^* \in \mathcal{R}_{\mathcal{K}}^{\mathcal{R}}$, $r_i^* = Rewrite(r_i)$and so hasPosIntent(t_1, t_2) $\notin Body(e)$ and hasNegIntent(t_1, t_2) $\notin Body(e)$ and $Head(e)$ is either of the form hasPosIntent(t_1, t_2) or hasNegIntent(t_1, t_2). Therefore for any two rules, $r_1^*, r_2^* \in \mathcal{R}_{\mathcal{K}}^{\mathcal{R}}$, $Atoms(Head(r_1^*) \cap Atoms(Body(r_2^*))$ = $\emptyset$. Then for any argument $e = \langle \{r_1^*, r_2^*\},\phi_2 \rangle$ s.t. $\phi_2 = Head(r_2^*)$, it must be the case that $\Gamma \cup r_2^* \vdash_{def} \phi_2$, and so there exists some other argument $e' = \langle \{r_2^*\}, \phi_2 \rangle$. However, since $\{r_2^*\} \subset \{r_1^*, r_2^*\}$ and $Claim(e) = Claim(e')$, e is not a valid argument (as it is not minimal). This holds for all e unless e is of size 1, when $Claim(\langle \{\}, \emptyset \rangle) \neq Claim(\langle \{e_2\} \phi_2 \rangle)$. □

So far I have presented quite general results about the relationships between arguments, but I now consider a special case. Recall that in general, I expect that we will author epistemic defeasible rules, and use a patient's values to rewrite them into ergonic ones. However, if we have an ergonic rule about a positive intention to take a treatment which is satisfied by the facts in the ontology, then we will develop an ergonic argument for that treatment, following which we will develop a committed argument for the same argument, which may then satisfy the body of the epistemic rule from which the ergonic was developed, in which case we can develop a hypothetical argument. The importance of this is that if we are arguing about treatments for a patient, this is a quite feasible chain of reasoning, as the proof below shows.

Proposition 5.4.20. *Let $r_c \in \mathcal{R}_{\mathcal{K}}^{\mathcal{C}}$ be a commital rule and $r \in \mathcal{R}_{\mathcal{K}}^{\mathcal{P}}$ be an epistemic rule, such that r_c, $r \in \Theta$ and $Head(r) \in \mathcal{V_P}^+$. Furthermore, let $\Gamma \subseteq \mathcal{G}_{\mathcal{K}}$ be a set of literals and let $\Delta \cup \Gamma \nvdash_{Ont} \bigwedge Body(r)$ and $\Delta \cup \Gamma \vdash_{Ont} \bigwedge Body(Rewrite(r))$. Then $(\Delta, \Theta) \vdash_{def} \phi$ and so there exists a hypothetical argument of the form $\langle \{Rewrite(r), r_c, r\}, \phi \rangle$.*

Proof. If $Head(r) \in \mathcal{V_P}^+$, then there is an ergonic rule $Rewrite(r) \in \mathcal{R}_{\mathcal{K}}^{\mathcal{R}}$ where $Head(Rewrite(r))$ = hasPosIntent(t_1,t_2). If $\Delta \cup \Gamma \vdash_{Ont} Body(Rewrite(r))$ there is an argument of the form $e = \langle \{Rewrite(r)\},\phi \rangle$ where ϕ = hasPosIntent(t_1,t_2). Then e is an ergonic argument in $\mathcal{A}_{\mathcal{R}}$, and $Commit(e) = \langle \{Rewrite(r),$ hasPosIntent(t_1,t_2) $\Rightarrow$hasTreatment(t_1,t_2)$\}$ hasTreatment(t_1,t_2)$\rangle$. Note that $Atoms(Body(r)) \setminus Atoms(Body(Rewrite(r))$ = {hasTreatment(t_1,t_2)}, and this is exactly the head of $Commit(e)$. Thus if $\Delta \cup \Gamma \vdash_{Ont} Body(Rewrite(r))$, $\Delta \cup \Gamma \cup$ {hasTreatment(t_1,t_2)} $\vdash_{Ont} Body(r)$ and so if $Head(r) = \phi$, $\Delta \cup \Gamma \cup$ {hasTreatment(t_1,t_2)} $\vdash_{Def} \phi$. We know that $\Delta \cup \Gamma \nvdash_{Ont} Body(r)$, and thus $\{Rewrite(r), h, r\}$ provides a minimal derivation for ϕ,and hence $\langle\{Rewrite(r), h, r\}, \phi\rangle$ is an argument. □

Indeed, the case study in chapter 7 is dominated by such arguments. Following on from 5.4.16 I can now prove the following proposition:

Proposition 5.4.21. *A purely ergonic argument is only present as a non-root node of a dialectical tree iff the argument it attacks is also an ergonic argument*

Proof. For an argument to appear as the non-root node in a tree, it must defeat its parent. An argument defeats another if it is a counter-argument to it and is preferred to it. However, from 5.4.16 above, if the parent is a pure epistemic argument, then a pure ergonic argument cannot be used as a counter argument to it, and therefore cannot defeat it. □

From this, and the earlier proof that pure ergonic arguments are only attacked by pure ergonic arguments, we can then see that if an OAF contains only purely ergonic and epistemic arguments, the dialectical trees contain arguments of one type only.

5.5 Comparison with other work

Mine is not the only work in argumentation to consider arguments for action, and I shall compare the work presented so far with the two other approaches that I reviewed in chapter 2. The reader may find it useful to revisit those pages, but I have summarised the notation here for ease of reading.

Amgoud There are three distinguished forms of rules:

1. Epistemic, of the form: $A_1 .... A_n \rightarrow \psi$
2. Recommending, of the form: $A_1 .... A_n \rightarrow d$
3. Decision, of the form: $A_1 .... A_n, d \rightarrow \psi$

Although arguments can be chained, there is no facility to rewrite rules of one form to another. In contrast, the paper by Bench - Capon & Prakken explicitly allows this.

Bench - Capon & Prakken

A rule is of the form: $\dfrac{\psi \Rightarrow \phi \;\; D\phi}{D\psi}$ where D is a modal operator, such that $D\phi$ means that ϕ is a positive goal.

5.5.0.1 Comparison to this work

I will start by noting some of the similarities between all three formalisms. Both of the existing formalisms and mine have defeasible rules in a logical language, some notion of desirable end-state or goal, a form of defeasible modus ponens and strong negation. These are unremarkable given the common intuitions underlying the work, but the differences are perhaps more revealing. Amgoud's work has a separate class of decisions, whereas in this thesis and Bench-Capon's, the decisions are made about elements of the language which may occur elsewhere in the rulebase. Furthermore, although Amgoud's framework captures the three classes of argument that my framework uses, with epistemic, recommending (c.f. ergonic) and decision (c.f. committed), she does not provide a mechanism for developing one form of argument from another, either at the level of the rules of the arguments. All formalisms share some ideas about defeasible rules, etc. Even where some are not shared (e.g. the use of ontology in OAF) there is a clear counterpart in the other frameworks (i.e. strict rules). However, there are some important areas where they differ and I will deal with these in turn.

Goals: All three formalisms include the notion of goals, and in all three they are formulae from the language. Bench-Capon & Prakken's formalism includes a single goal base, G, which may contain either positive (Dx) or negative ($\neg Dx$) elements, whereas both Amgoud's work and mine contain two goal bases, one of positive and negative goals. In all cases the set of goals is assumed to be consistent.

Decisions: All the frameworks have the idea that certain elements of the language may be considered as the subject of decisions and all three apply this distinction across elements of the language (by denoting them 'choosable' or elements of d, or being treatments respectively). However, the choosable formulae in BCP are defined at the level of the language, whereas mine are in a function external to the language. I would argue that there are occasions when an option is choosable, and others when it may not be, and this flexibility cannot be encoded in formulae at the level of the language. Thus implementing a more flexible rewriting system (such as conditional rewriting, based both some preconditions but only applying to certain atoms) would be easier in OAF.

Generation of Decisions: All three systems have notions of goals; all three have notions of 'choosable' elements of the languages in order to achieve them. However, Amgoud's work contains no mechanism for generating arguments about actions from rules about the effects of those actions, as she has no mechanism for re-writing rules, unlike both of the other approaches. Her work therefore fails to make one of the key inferences of the (informal) practical syllogism, as it has no way of reasoning from information about the effects of an action to an argument for that action. Such differences become increasingly important the larger the knowledge-base becomes, where the cost of rewriting rules becomes prohibitive. In contrast, both BCP and OAF allow one to generate a new ergonic rules from a set of epistemic rules and goals (values). However, BCP allows one to continue to process through several iterations, whereas OAF does not. On the one hand, this allows BCP to be more succinct, as a small set of goals can generate a great deal of rewritten rules. However, from a practical point of view, this has some problems, as it seems difficult to predict the exact consequences of altering the goal base on a set of rules. In contrast, rewriting in OAF is a single step process, and so checking the relationship between goals and rules is relatively simple. Furthermore, although all three approaches start with a consistent set of goals, BCP gives no guarantee that this is preserved, as the following example shows.

Example 5.5.1. Given the following rules:

$$a \wedge b \Rightarrow c$$
$$d \wedge b \Rightarrow e$$

and the goal-base

G= $\{Dc, \neg De\}$

with b being choosable, then we can form an argument for Db

$$\frac{a \wedge b \Rightarrow c \;\; Dc}{Db}$$

but also an argument for $\neg Db$

$$\frac{d \wedge b \Rightarrow e \;\; \neg De}{\neg Db}$$

As a result the new set of goals G is $\{Dc, \neg De, Db, \neg Db\}$ which is inconsistent. This is because BCP allows the inference of new goals, whereas OAF and the work by Amgoud do not.

Intentions, Actions and Goals: All three frameworks have some notion of an argument for an action; however, only OAF explicitly distinguished between a goal, an action and an intention. In BCP, the backward chaining aspect of the modal operator D seems attractive, but is based on there being no distinction between a desired end state (a goal) and the intention to carry out some action to achieve it. While one might be happy to loosely talk about having a 'goal' of achieving some action, to conflate the that with a goal of some end-state seems to mix process and outcome to an uncomfortable degree. Furthermore, the restriction of the use of D to a single propositional literal makes this backward chaining element rather fragile - it is not hard to consider some outcome that requires a conjunction of actions. Amgoud's work does not suffer from this problem, but instead makes no link between an argument for a decision and that decision having been carried out. Thus, if the set of decisions is disjoint from the other elements of the language, then a rule about the effects of some decision will never be satisfied by an argument for that decision; this is why she has a separate class of decision arguments, which contain a decision in their body and argue for some outcome. Again, this achieves that same as the work in OAF on committed reasoning, but only by the addition of another set of rules which link decisions with their outcomes. In general, therefore, I would argue that the distinction between goals and intentions is too weak (i.e. non-existent) in one, and the link between decisions and their outcomes too weak in the other, and both weakness cause substantial problems.

Negated Goals: BCP explicitly uses the relationship $\neg D \neg \phi \equiv D\phi$ to related negated literals and goals. As I noted earlier, we do not do this in OAF, as ¬Men(x) is the set of all things that are not members of the set Men(x).

Committed Reasoning: Amgoud's framework allows one to reason about the outcome of some possible decision, by using a separate class of decision arguments. Bench-Capon's work does not allow this, as it allows no mechanism for reasoning from the intention to carry out an action to its consequences. This is necessary because the lack of distinction between goals and intentions means that the 'commital' of an intention (that is, reasoning from Dx to x) would also result in the commital of a goal. This is comes close to formalising 'wishful thinking' - the desire that something be so makes it so, and while useful for reasoning about actions, is difficult when reasoning about goals. Although Amgoud's work does make this distinction, it is at the expense of having another class of rule that explicitly refers to decisions in the body of the rule. In contrast, the commital mechanism of OAF allows one to reuse the same information automatically.

Preferring arguments based on type: Amgoud makes an explicit assertion that for any epistemic argument A_e and recommending argument A_r and decision argument A_d, then A_e is preferred to A_r which is preferred to A_d. This is defended on the grounds that epistemic and recommending arguments are to be believed more than decision ones, but epistemic ones are always to be preferred to recommending ones. From this, she derives an asymmetric defeat relationship between epistemic and non-epistemic arguments, based on this strict hierarchy of rules. In contrast, I have demonstrated a similar asymmetry in attack relations (and hence defeat) but based entirely on the elements of the rules. I would suggest that this is preferable, as the restrictions on the atoms in the rules are small, and help delineate the different types of rules; the relationships between arguments come from this, rather than preferences between rules. The problem is that it does not seem hard to conceive of a situation when a recommending or decision argument is much more convincing than an epistemic one, as 5.5.2 demonstrates.

Example 5.5.2. Weak evidence for Meningitis.

Consider a propositional language $\mathcal{L}$ with the symbols k, n, m, a, s, p and the following interpretations:

k: Kerning's sign positive
n: Neck stiffness present
m: Bacterial meningitis
a : Give IV antibiotics
s : Seriously ill
p : Purpuric rash present

rules of the form:

$s \Rightarrow a$
$p \Rightarrow m$
$k \Rightarrow m$
$\neg n \Rightarrow \neg m$
$m \Rightarrow a$

and a set of facts, $\Gamma = \{s, p, k, \neg n\}$ to represent a patient who is seriously ill, has a purpuric rash and Kerning's sign positive but no neck stiffness, then we can construct the arguments:

$\langle \{s \Rightarrow a\}, a \rangle$
$\langle \{p \Rightarrow m\}, m \rangle$
$\langle \{k \Rightarrow m\}, m \rangle$
$\langle \{\neg n \Rightarrow \neg m\}, \neg m \rangle$
$\langle \{k \Rightarrow m, m \Rightarrow a\}, a \rangle$
$\langle \{p \Rightarrow m, m \Rightarrow a\}, a \rangle$

We have a situation where we have a patient who is seriously ill (s) and has some symptoms suggestive of meningitis (p, k) but some that are against it $(\neg n)$. We can construct an argument for giving antibiotics based on the patient's condition, and can also construct another argument that the patient may have meningitis, and therefore should have antibiotics. However, the arguments that he may have meningitis are attacked, whereas the argument that he should have some antibiotics is not; therefore (ignoring for now which approach we take to resolving argument conflict), in contrast to Amgoud's approach it seems that our ergonic argument should be considered to be 'stronger' than than either of the arguments for believing that the patient has meningitis.

5.6 Summary

This chapter has introduced the idea of committed arguments, and provided definitions of such arguments as well as a procedure for generating them. In addition, we have seen the importance of constraint rules in generating conflict between our arguments. However, in order to resolve our arguments in dialectical trees, we need to define preferences between our arguments, and it is this that I now turn to.

Chapter 6

Preferences & Implementation

So far, this thesis has concentrated on defining an ontology based argumentation formalism (OAF), and then extending this to capture aspects of ergonic and hypothetical reasoning. In the next chapter, I present a substantial case study to assess the work, but before I do this, I need to address two subsidiary issues. Firstly, I show how OAF resolves conflict between arguments, using preferences between rules to decide on preference between arguments. This is similar to the original proposal by Garcia & Simari, but I extend an existing clinical evidence ranking scale to determine a simple ordering on rules, and I use this to provide a concrete proposal for the case study.

Secondly, the size of the case study makes it difficult to continue with the current manual approach. I therefore present an outline of a prototype implementation that adapts an existing rule-engine for argument generation. The prototype implements argument generation using a novel algorithm, and I give some simple proofs of its soundness, completeness and an upper-bound on its time-complexity.

6.1 Preferences

In Chapter 3, I gave a definition of argument comparison based on the priorities of the rules in the support. In the original DeLP paper [32], Garcia & Simari suggest that we abstract from any exact mechanism and merely assume that we have some way of deciding whether one argument defeats another. However, for the case study in the next chapter, I need a concrete proposal for deciding on defeat status. The first step is to describe a simple ordering on the elements of a set of rules. I refer to the place of a rule in this order as the ***rank*** of a rule.

Definition 6.1.1. Let $\mathcal{K}$ be some vocabulary and $\mathcal{R}_{\mathcal{K}}$ the set of defeasible rules. Let $\Theta \subseteq \mathcal{R}_{\mathcal{K}}$ be some set of defeasible rules, such that the elements of Θ are ranked. Then a function $Rank : \mathcal{R}_{\mathcal{K}} \mapsto \mathbb{N}^{+}$ returns the rank of a rule $r \in \Theta$.

The ranking of a set of rules is an extra piece of knowledge that we associate with an OAF. I shall shortly give a simple way of representing the ranking of rules.

Example 6.1.2. Let Θ = {NICE2003, NLCN1997, BMJ2001} where the rules are as follows:

NICE2003: People(x)

$\wedge$ hasDisease(x,y) $\wedge$ BreastCancerTypes(y)

∧ Tamoxifen5YrCourse(z)

⇒ hasTreatment(x,z)

NLCN1997: People(x)

∧ hasDisease(x,y) ∧ BreastCancerTypes(y)

∧ Tamoxifen2YrCourse(z)

⇒ hasTreatment(x,z)

BMJ2001: People(x)

∧ hasDisease(x,y) ∧ BreastCancerTypes(y)

∧ ECMFChemotherapy(z)

⇒ hasTreatment(x,z)

such that $Rank$(NICE2003) = 1, $Rank$(NLCN1997) = 2, $Rank$(BMJ2001) = 3

Since the support of an argument is a set of rules, we need some way of extending our ranking from single rules to sets of rules. The original DeLP paper defines a system for checking the preference status of arguments considering all the rules in the support. However, if we just consider two arguments, the support of each argument can contain a rule that is preferred to a rule in the other argument, and thus neither argument is preferred to the other. Also, every rule in each argument needs to be compared to every rule in every other argument. Instead, I calculate the average rank of the set of rules in the support. This has the advantage of being simple to calculate, and for each rule the value can be calculated once, and the average rank cached and simply looked-up. I therefore use it for the case study in the next chapter.

Definition 6.1.3. Let $\mathcal{K}$ be some vocabulary and $\mathcal{R}_\mathcal{K}$ the set of defeasible rules. Let $\Theta \subseteq \mathcal{R}_\mathcal{K}$ be some set of defeasible rules, and $Rank$ is a ranking function for Θ. Then the function $AvRank$: $2^{\mathcal{R}_\mathcal{K}} \mapsto \mathbb{R}^+$ returns the average rank of the rules in Θ, such that if $\Theta = r_1...r_n$, $AvRank(\Theta) = \frac{Rank(r_1)+...+Rank(r_n)}{|\Theta|}$

Example 6.1.4. Let Θ = {NICE2003, NLCN1997, BMJ2001} such that $Rank$(NICE2003) = 1, $Rank$(NLCN1997) = 2, $Rank$(BMJ2001) = 3. Then $AvRank$({NICE2003, NLCN1997}) = 1.5

Given two arguments that conflict, I use the rankings to decide which argument is preferred. We do this by comparing the average ranking of the support of the two arguments.

Definition 6.1.5. Let $\mathcal{K}$ be some vocabulary and $\mathcal{R}_\mathcal{K}$ be a set of defeasible rules. Let $\Theta \subseteq \mathcal{R}_\mathcal{K}$ be some set of ranked defeasible rules. Given two argument structures $\langle A_1, \phi_1 \rangle$ and $\langle A_2, \phi_2 \rangle$, the argument $\langle A_1, \phi_1 \rangle$ will be preferred over $\langle A_2, \phi_2 \rangle$ iff $AvRank(A_1) > AvRank(A_2)$

What remains unanswered in both this and the original definition is from *where* this preference relationship comes.

6.1.1 A Clinical Approach

As described in Chapter 1, we already know that the results of clinical trials can conflict. In some cases, we resolve this by considering whether there is a difference in the quality of the trials. However, any

Level of Evidence	Source	Comments
1a	Systematic Review (SR) of RCTs	Low heterogeneity
1b	Single RCT	
1c	'All-or-None' Case Series	E.g. Used for evaluating diagnostic tests
2a	SR of Cohort studies/ 2b studies	
2b	Cohort Study or poor RCT	E.g. <80% follow-up
2c	'Outcomes' Research	Observational study; Natural experiments
3a	SR of Case-Control study	
3b	Individual Case-Control study	
4	Case Series	
5	Expert Opinion, Basic Science, Physiology	

Table 6.1: Grading of Evidence for Clinical Medicine. **Source** refers to the type of study that was done. **Level of Evidence** is a summary label to make it easier to annotate guidelines and discussions of the evidence supporting statements. Note that the level of evidence considers different factors such as the type of study, the completeness of the data, etc. and summarises them in a single measure. A fuller explanation is given in [15].

such system needs to strike a middle ground between completeness and ease of use. The current system for grading clinical evidence is given in Table 6.1. This relies mainly on the type of study, and how well it was performed, to grade the evidence. Although it delivers a concise summary of the evidential level that supports a statement, for many purposes it is too coarse. To improve our discrimination between papers, we need to consider more than just the level of evidence. The range of information we could consider is almost limitless, and here I give only one possible approach, rather than a definitive solution.

6.1.2 Refining Clinical Preferences

For the purpose of the thesis, I use information about the source on which the rule is based. Consider a medical study which we can categorise in many different ways - by year of publication, author affiliation, journal of publication, etc. It is useful to subdivide the information that we have about a study, and I assume that there are domain-specific sets that represent fields such as year of publication, etc. I shall refer to such a set as a ***criterion***. Each defeasible rule has some label, E, which we use to associate rules with information about their sources, such that every E is associated with a single element of each ***criterion***.

Example 6.1.6. Consider a clinical trial about the effects of Tamoxifen. The study recruits 7325 patients, all from the UK, and is published in the BMJ in 2004. I associate a label $\mathsf{BMJ2004}$ with a rule developed from the study and consider a criterion such as the year of publication (denoted Y), where the elements of Y are the integers 1900 - 2010. I associate the label, and hence the rule, with the element $2004 \in Y$.

In the rest of the thesis, I shall assume that every label is associated with four criteria.

Definition 6.1.7. Let Y, N, L, A be four sets. Then Y and N are each a set of positive integers (i.e.

$\mathbb{Z}^+$) representing the year of publication and the number of patients in the study. L is the set of levels of evidence as in the table above and A is a set of geographical areas where A = {World, Europe, UK, USA, Other}.

The intention is then to use the information in each criterion to rank the labels. I assume that we can define a total ordering over the elements of each criterion, and I shall assume that each criterion is ordered as follows:

L : $1a$ is preferred to $1b$ which is preferred to $1c$, and so on until 5, which is the least preferred

N : Larger numbers (of patients) are preferred to smaller numbers

Y : Larger numbers (representing more recent years) are preferred

A : $World$ is preferred to $Europe$ which is preferred to UK which is preferred to USA which is preferred to $Other$

I have presented a method for ranking rules with respect to a criterion. However, we need to know how to use multiple criteria, which we do by considering one criterion, and only considering the next if the rules are considered equal when judged on that.

Definition 6.1.8. A ***precedence order*** is a finite sequence of criteria, denoted by the letters of the criteria ordered in the sequence in which they take precedence. Given a precedence ordering R = $[X_1,\ldots, X_n]$ we obtain $Rank_R$ as follows: for a pair of rules r_1, $r_2 \in \mathcal{R}_\mathcal{K}$, $Rank_R(r_1)$ > $Rank_R(r_2)$ when r_1 is preferred to r_2 according to a criterion X_i and for all $j > i$, r_1 and r_2 are equal according to criteria X_j. $Rank_R(r_1)$ = $Rank_R(r_2)$ iff r_1 and r_2 are equally preferred according to criterion X_i for $1 \leq i \leq n$.

Example 6.1.9. Consider the domain ordering that first considers the level of evidence (L) and then the number of patients involved (N), which we denote $[L, N]$. Then for two rules r_1 and r_2, we first look at the level of evidence for r_1 and r_2. If it is different, the rule with the higher level is preferred. If it is equal, then we proceed to look at the number of patients involved in the trial; whichever has the greater is the preferred rule. In contrast, the ordering based on the year of publication (Y) and the geographical area (A), denoted [Y, A], would consider first the year of publication and then the area.

I have given a definition to allow us to calculate preferences between sets of rules, and a way of specifying the basis on which we calculate the preferences. I now extend the definition of an OAF to include this.

Definition 6.1.10. Let $\mathcal{K}$ be a vocabulary, $\mathcal{O}_\mathcal{K}$ the set of ontology formulae, $\mathcal{R}_\mathcal{K}$ the set of defeasible rules and $\mathcal{G}_\mathcal{K}$ the set of literals. Also, let p be in $PeopleNames(\Omega)$, $\mathcal{V}_P^+ \subseteq \mathcal{G}_\mathcal{K}$ be the set of positive values for patient p and $\mathcal{V}_p^- \subseteq \mathcal{G}_\mathcal{K}$ the set of negative values for patient p. Let V_P = ($\mathcal{V}_p^+$,$\mathcal{V}_p^-$) be the patients values. A ***Ranked Ontology-based Argumentation Framework*** (R-OAF), denoted Ω, is a tuple (Θ, Δ, V_p, R) where $\Theta \subseteq \mathcal{R}_\mathcal{K}$ is a finite set of defeasible rules, $\Delta \subseteq \mathcal{O}_\mathcal{K}$ is an ontology, V_p = ($\mathcal{V}_p^+$, $\mathcal{V}_p^-$) and R is a precedence order.

The main purpose of preferences is to allow us to decide on defeat status between conflicting arguments. The example below extends the one given in the last two chapters. It defines a preference ordering on the arguments based on the $[L, Y, A]$ ordering and uses this to calculate the defeat status of arguments, and hence allow us to construct dialectical trees.

Example 6.1.11. I repeat the previous example, but add information on preferences and the sources of the rules.

Let $\mathcal{K}$ be a vocabulary, $\mathcal{G}_{\mathcal{K}}$ the set of literals, $\Gamma \subseteq \mathcal{G}_{\mathcal{K}}$ a set of facts, Θ a set of epistemic defeasible rules and Δ_{BC} the breast cancer ontology. Let $V_{MsJones}$ be the sets of positive and negative values for MsJones, and Φ be set of ergonic defeasible rules. Then Ω_{Δ} = $(\Theta, \Gamma \cup \Delta_{BC}, V_{MsJones}, R)$ is an R-OAF, where Θ, Φ, Γ, $V_{MsJones}$ and R are as below.

R = $[L, Y, A]$

We have the same arguments as before, which I repeat to aid the discussion:

Γ = {Woman(MsJones), hasDisease(MsJones,ProtoBreastCancer)}

$\mathcal{V}^{+}_{MsJones}$ ={hasDeltaRisk(MsJones,IncreasedOS1.01),

hasDeltaRisk(MsJones,IncreasedOS1.02),

hasDeltaRisk(MsJones,IncreasedOS1.03),

hasDeltaRisk(MsJones,IncreasedOS1.1),

hasDeltaRisk(MsJones,IncreasedOS1.2)}

and

$\mathcal{V}^{-}_{MsJones}$ ={hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer2.3),

hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer2.53),

hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer4.1),

hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer6.0),

hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer6.4),

hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer7.5)

hasDeltaRisk(MsJones, DecreasedOS0.7)}.

$V_{MsJones}$ = $(\mathcal{V}^{+}_{MsJones}, \mathcal{V}^{-}_{MsJones})$

Θ = {NICE2003, NLCN1997, NCRN2001, BMJ1999, BJC2004}

where

NICE2003: : (People(x)

$\wedge$ hasDisease(x,y) $\wedge$ BreastCancerTypes(y)

$\wedge$ Tamoxifen5YrCourse(z)

$\Rightarrow$ hasTreatment(x,z)

NLCN1997: : (People(x)

$\wedge$ hasDisease(x,y) $\wedge$ BreastCancerTypes(y)

$\wedge$ Tamoxifen2YrCourse(z)

$\Rightarrow$ hasTreatment(x,z)

NCRN2001: : (People(x)

$\wedge$ hasDisease(x,y) $\wedge$ BreastCancerTypes(y)

$\wedge$ ECMFChemotherapy(z)

$\Rightarrow$ hasTreatment(x,z)

BMJ1999: : (People(x)

$\wedge$ hasTreatment(x,y) $\wedge$ TamoxifenByLength(y)

$\wedge$ IncreasedOS(z) $\wedge$ hasValue(z,1.2)

$\Rightarrow$ hasDeltaRisk(x,z)

BJC2004: : (People(x)

$\wedge$ hasTreatment(x,y) $\wedge$ TamoxifenByLength(y)

$\wedge$ IncreasedRiskEndometrialCancer(z) $\wedge$ hasValue(z,4.1)

$\Rightarrow$ hasDeltaRisk(x,z)

and

Φ = $\{Rewrite(r)$ | $r \in \Theta$ and $Rewrite(\Theta) \neq null\}$

and we have the following commital rules:

r_{c1}^{+} : hasPosIntent(MsJones, TamD2YrCourse) $\Rightarrow$ hasTreatment(MsJones, TamD2YrCourse)

r_{c2}^{+} : hasPosIntent(MsJones, TamE5YrCourse) $\Rightarrow$ hasTreatment(MsJones, TamE5YrCourse)

r_{c3}^{+} : hasPosIntent(MsJones, TamFgt5YrCourse) $\Rightarrow$ hasTreatment(MsJones, TamFgt5YrCourse)

r_{c1}^{-} : hasNegIntent(MsJones, TamD2YrCourse) $\Rightarrow$ $\neg$hasTreatment(MsJones, TamD2YrCourse)

r_{c2}^{-} : hasNegIntent(MsJones, TamE5YrCourse) $\Rightarrow$ $\neg$hasTreatment(MsJones, TamE5YrCourse)

r_{c3}^{-} : hasNegIntent(MsJones, TamFgt5YrCourse) $\Rightarrow$ $\neg$hasTreatment(MsJones, TamFgt5YrCourse)

Heuristic Arguments, $\mathcal{A}_{\mathcal{K}}^{\mathcal{U}}$:

a_1 $\langle$\{NICE2003\}, hasTreatment(MsJones,TamE5YrCourse)$\rangle$

a_2 $\langle$\{NLCN1997\}, hasTreatment(MsJones,TamD2YrCourse)$\rangle$

a_3 $\langle$\{NCRN2001\}, hasTreatment(MsJones,ECMFTypeInstance)$\rangle$

Ergonic Arguments, $\mathcal{A}_{\mathcal{K}}^{\mathcal{R}}$:

a_4 $\langle$\{BMJ1999$^*_1$\}, hasPosIntent(MsJones,TamD2YrCourse)$\rangle$

a_5 $\langle$\{BMJ1999$^*_2$\}, hasPosIntent(MsJones,TamE5YrCourse)$\rangle$

a_6 $\langle$\{BMJ1999$^*_3$\}, hasPosIntent(MsJones,TamFgt5YrCourse)$\rangle$

a_7 $\langle$\{BJC2004$^*_1$\}, hasNegIntent(MsJones,TamD2YrCourse)$\rangle$

a_8 $\langle$\{BJC2004$^*_2$\}, hasNegIntent(MsJones,TamE5YrCourse)$\rangle$

Epistemic Arguments, $\mathcal{A}_{\mathcal{K}}^{\mathcal{P}}$:

a_9 $\langle$\{BJC2004$^*_3$\}, hasNegIntent(MsJones,TamFgt5YrCourse)$\rangle$

a_{10} $\langle$\{NICE2003*, BMJ1999\}, hasDeltaRisk(MsJones, IncreasedOS1.1)$\rangle$

a_{11} $\langle$\{NLCN1997*, BMJ1999\}, hasDeltaRisk(MsJones, IncreasedOS1.1)$\rangle$

a_{12} $\langle$\{NICE2003*, BJC2004\}, hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer2.3)$\rangle$

a_{13} $\langle$\{NLCN1997*, BJC2004\}, hasDeltaRisk(MsJones, IncreasedRiskEndometrial-Cancer2.3)$\rangle$

Committed Arguments, $\mathcal{A}_{\mathcal{K}}^{\mathcal{C}}$:

a_{14} $\langle$\{BMJ1999$^*_1$, r^+_{c1}\}, hasTreatment(MsJones,TamD2YrCourse)$\rangle$

a_{15} $\langle$\{BMJ1999$^*_2$, r^+_{c2}\}, hasTreatment(MsJones,TamE5YrCourse)$\rangle$

a_{16} $\langle$\{BMJ1999$^*_3$, r^+_{c3}\}, hasTreatment(MsJones,TamFgt5YrCourse)$\rangle$

a_{17} $\langle$\{BJC2004$^*_1$, r^-_{c1}\}, ¬hasTreatment(MsJones,TamD2YrCourse)$\rangle$

a_{18} $\langle$\{BJC2004$^*_2$, r^-_{c2}\}, ¬hasTreatment(MsJones,TamE5YrCourse)$\rangle$

a_{19} $\langle$\{BJC2004$^*_3$, r^-_{c3}\}, ¬hasTreatment(MsJones,TamFgt5YrCourse)$\rangle$

Hypothetical Arguments, $\mathcal{A}_{\mathcal{K}}^{\mathcal{H}}$:

a_{20} $\langle$\{BMJ1999$^*_1$, r^+_{c1}, BMJ1999\}, hasDeltaRisk(MsJones, IncreasedOS1.1)$\rangle$

a_{21} $\langle$\{BMJ1999$^*_2$, r^+_{c2}, BMJ1999\}, hasDeltaRisk(MsJones, IncreasedOS1.1)$\rangle$

a_{22} $\langle$\{BMJ1999$^*_3$, r^+_{c3}, BMJ1999\}, hasDeltaRisk(MsJones, IncreasedOS1.1)$\rangle$

a_{23} $\langle$\{BMJ1999$^*_1$, r^+_{c1}, BMJ1999\}, hasDeltaRisk(MsJones, IncreasedRiskEndometrial-Cancer2.3)$\rangle$

a_{24} $\langle$\{BMJ1999$^*_2$, r^+_{c2}, BMJ1999\}, hasDeltaRisk(MsJones, IncreasedRiskEndometrial-Cancer2.3)$\rangle$

a_{25} $\langle$\{BMJ1999$^*_3$, r^+_{c3}, BMJ1999\}, hasDeltaRisk(MsJones, IncreasedRiskEndometrial-Cancer2.3)$\rangle$

Source	Year	Evidence Level	Area
NICE2003	2003	2c	USA
NLCN1997	1997	2c	UK
NCRN2001	2001	2b	Europe
BMJ1999	1999	1a	Europe
BJC2004	2004	1b	UK

Table 6.2: Information on the sources of the rules. ***Source*** is the label of the rules; ***Year*** is the year of publication; ***Evidence Level*** refers to the clinical evidence level as described earlier; ***Area*** refers to the geographical area from where the patients are drawn.

I use the precedence order $[L, Y, A]$, on the basis of which we can see that

BMJ1999 > BJC2004 > NCRN2001 > NICE2003 > NLCN1997

The difference is that now we are in a position to determine the preferences between rules, and hence arguments, and so construct our dialectical trees, as in chapter 3. Recall that we start by choosing a root node, which in this case is a_1: $\langle$\{NICE2003\}, hasTreatment(MsJones,TamE5YrCourse)$\rangle$, an argument that MsJones is having treatment with a 5 Year course of tamoxifen. The leaves of the tree are the arguments that defeat the root node, and they in turn have leaves that are arguments that defeat them. I label the nodes of the tree with the label of the corresponding argument.

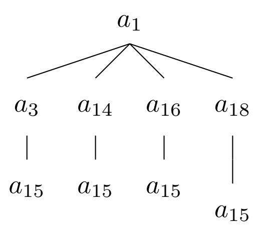

The root node is a_1: $\langle$\{NICE2003\}, hasTreatment(MsJones,TamE5YrCourse)$\rangle$. Those defeating it are :

a_3 $\langle$\{NCRN2001\}, hasTreatment(MsJones,ECMFTypeInstance)$\rangle$

a_{14} $\langle$\{BMJ1999$^*_1$, r^+_{c1}\}, hasTreatment(MsJones,TamD2YrCourse)$\rangle$

a_{16} $\langle$\{BMJ1999$^*_3$, r^+_{c3}\}, hasTreatment(MsJones,TamFgt5YrCourse)$\rangle$

a_{18} $\langle$\{BJC2004$^*_2$, r^-_{c2}\}, ¬hasTreatment(MsJones,TamE5YrCourse)$\rangle$

Of interest, only one of these is an attack because of syntactic negation - the others attack a_1 because they suggest a different treatment; also 5 of the 6 arguments are committed, and committed arguments have ergonic subarguments (consistent with proposition 5.4.6). Most notably, the root node, a_1 is warranted, but it is only warranted because of the presence of another argument with the same claim and a higher priority, a_{15} $\langle$\{BMJ1999$^*_2$, r^+_{c2}\}, hasTreatment(MsJones,TamE5YrCourse)$\rangle$. The final observation is that because we have not defined any constraint rule that links intention and action, arguments for an intention to do x not directly attack those for an argument that y is underway. This is an example of the proof at the end of chapter 5, where I showed that such a attack could not happen; here, we have a tree with a non-ergonic root node, and ergonic arguments are only involved as counter-arguments to hypothetical one that argue (against the root) for a different action.

6.1.3 Summary

In this half of the chapter, I have presented a simple domain specific technique for describing a preference relationship between rules and sets of rules, and have shown how this can be used to determine the defeat status of arguments in a dialectical tree. The approach I have taken here should be considered as an example of how we might resolve such issues in a real situation, rather than as a definitive answer. Specifically, we might want to consider more criteria when ranking the rules, such as journal of publication, etc. One of the implications of this ranking scheme is that seemingly minor changes in the ordering of criteria, or which criteria we choose to use can have considerable impact on the eventual preference ordering of rules. For example, the use of the number of patients in a study as a criterion gives a very fine-grained ordering on rules, as the number of patients in trials varies, and is almost never the same in two different trials.

6.2 Implementation

In the next chapter I present a substantial worked example. Developing the arguments manually is time consuming and error-prone and I have therefore produced a minimal prototype, guided by a few principles. The first of these is that, as far as possible, it should reuse existing software components; since the work is based on the $\mathcal{SROIQ}(\mathcal{D})$ description logic, which is the basis for OWL 2, I can reuse existing tools. Secondly, I have opted to focus on argument generation, as that is the novel aspect of my thesis - dialectical tree construction is essentially unchanged from previous work. Given these constraints, I have a choice of various software tools, none of which do all of what I need. Given the choices, I have opted to use an existing RDF/OWL platform, which includes a rule-engine, based the prototype on the Jena framework and Pellet, a leading OWL reasoner. Both are open-source and Java based, and Jena provides an API for reading and manipulating RDF (and hence OWL) files, as well as a simple RDF-based rule engine. Pellet is a sound and complete tableau reasoner for $\mathcal{SROIQ}(\mathcal{D})$, and can use Jena data structures.

In developing the algorithms below, therefore, I make some assumptions. Specifically, I shall assume that we have access to a rule engine that when supplied with an ontology and a set of rules, returns the *set* of literals that are inferred by the rules. However, the rule engine does not do ontological reasoning, and so it does not return literals which are ontologically entailed. This is a reasonable assumption as it describes some of the functionality of the Jena package, and is consistent with the claims of our foreground arguments, but we need to note that returning a set of literals is slightly different to our existing definition of defeasible entailment, and I will therefore use a function to refer to it. In the definitions below I use $\wp(s)$ to denote the powerset of a set s.

Definition 6.2.1. Let $\mathcal{K}$ be a vocabulary, Δ be a DL ontology, $\mathcal{G}_\mathcal{K}$ be the set of literals, $\mathcal{R}_\mathcal{K}$ the defeasible rule language and $\Theta \subseteq \mathcal{R}_\mathcal{K}$ a set of defeasible rules. Then the ***rule engine*** function $Inf : \wp(\mathcal{R}_\mathcal{K}) \times \wp(\mathcal{O}_\mathcal{K})$

$\mapsto \wp(\mathcal{G}_{\mathcal{K}})$ returns the set of all literals $\{\phi_1 ... \phi_n\}$ such that for each $\phi_i \in \{\phi_1 ... \phi_n\}$:

1. $(\Theta, \Delta) \vdash_{Def} \phi_i$
2. $\Delta \nvdash_{Ont} \phi_i$
3. $\phi_i = Head(r)$ for some $r \in \Theta$

Note that Inf returns literals that have a defeasible derivation but do not have a strict derivation. This definition is consistent with literals that are the claims of foreground arguments, although it is stricter than that of defeasible derivation given in chapter 3.

The other consideration is that in developing arguments from a set of rules Θ one method would be to pick elements of $\wp(\Theta)$. Many elements of Θ cannot be used as the support of an argument as they are either inconsistent or non-minimal. However, dealing with these two 'violations' requires different approaches. Given a set of rules (i.e. an element of $\wp(\Theta)$), one can check whether it is consistent by checking that the union of facts and inferences is consistent. In the remainder of the work, I shall concentrate on ensuring consistency, not minimality, and confine the rest of this discussion to the generation of non-minimal foreground (NMF) arguments, defined below.

Definition 6.2.2. Let Δ be some DL ontology, $\mathcal{G}_{\mathcal{K}}$ the set of all literals and $\mathcal{R}_{\mathcal{K}}$ the defeasible rule language. For some literal $\phi_i \in \mathcal{G}_{\mathcal{K}}$ and a set of rules $A \subseteq \mathcal{R}_{\mathcal{K}}$, then $\langle A, \phi_i \rangle$ is a ***NMF argument*** for ϕ_i, iff:

1. $(A, \Delta) \vdash_{Def} \phi_i$
2. $\Delta \nvdash_{Ont} \phi_i$
3. $\phi_i = Head(r)$ for some $r \in A$
4. A is not contradictory

Note the close similarity between the claims of NMF arguments and the literals returned by Inf; the conditions on them are the same, except that NMF arguments also requires the set of rules to be non-contradictory.

Given the definition of a non-minimal argument, it will be helpful to have a function that returns the set of all non-minimal arguments from an OAF.

Definition 6.2.3. Let $\mathcal{K}$ be some vocabulary, $\mathcal{O}_{\mathcal{K}}$ be the set of ontological formulae, $\Delta \subseteq \mathcal{O}_{\mathcal{K}}$ be some DL ontology, $\mathcal{G}_{\mathcal{K}}$ the set of all literals, $\mathcal{R}_{\mathcal{K}}$ the set of defeasible rules and $\mathcal{A}_{\mathcal{K}}$ the set of arguments. Let $\Theta \subseteq \mathcal{R}_{\mathcal{K}}$ be some set of defeasible rules. The function $NMFArgs$: $\wp(\mathcal{R}_{\mathcal{K}}) \times \wp(\mathcal{O}_{\mathcal{K}}) \mapsto \wp(\mathcal{A}_{\mathcal{K}})$ returns the set of NMF arguments from (Θ, Δ).

This chapter contains two algorithms to generate arguments given an ontology and a set of rules. It is important to ensure that such algorithms are correct, in the sense that they return only the arguments that they should, and return all of the arguments that they should. However, I need to be clear about what I mean when by correctness. For the purposes of this chapter, I make two assumptions: Firstly, given some ontology Δ and set of rules Θ, I assume that the support of every argument is an element of the powerset of Θ. Therefore, we can assess an algorithm for argument generation against the output of $NMFArgs$, and use this as a guide as to whether an algorithm is sound and complete.

Definition 6.2.4. Let $\mathcal{K}$ be a vocabulary, Δ be a DL ontology, $\mathcal{R}_{\mathcal{K}}$ the defeasible rule language and $\Theta \subseteq \mathcal{R}_{\mathcal{K}}$ a set of defeasible rules. Let $Alg : (\Theta, \Delta) \mapsto \wp(\mathcal{A}_{\mathcal{K}})$ be some function for argument generation that returns a set of arguments. Then Alg is ***sound*** with respect to $NMFArgs$ iff for every element in the output of $Alg(\Theta, \Delta)$ is in the output of $NMFArgs(\Theta, \Delta)$ and is ***complete*** with respect to $NMFArgs$ iff for every $a \in NMFArgs(\Theta, \Delta)$, $a \in Alg(\Theta, \Delta)$.

6.3 A Simple Approach

I shall start by assuming that we have a set of ground rules. The simplest approach to constructing arguments is to generate the power set of the rules and use each element of the powerset in turn. The pseudocode for this approach is given in below by MAKE-PSARGS, where I make use of the $Literals_{\mathcal{K}}(\Delta)$ function from chapter 2 that returns the set of all literals from an ontology.

MAKE-PSARGS(Θ, Δ)

1 $outArgs \leftarrow \{\}$
2 **for** $element \in \wp(\Theta)$
3 **if** $element$!= { }
4     $results \leftarrow Inf(element, \Delta)$
5   **for** $\phi \in results$
6   **if** $conflict_{\Delta}(Literals_{\mathcal{K}}(\Delta), \phi)$ is false
7             $outArgs \leftarrow outArgs \cup \{\langle\, element, \phi \,\rangle\}$
8 **return**$(outArgs)$

Since I will reuse many of the ideas in this algorithm again, it is worth discussing. On line 2, we pick an element of $\wp(\Theta)$, check that it is not the empty element (line 3) and use it with the set of facts to infer a set of formulae (line 4). We use these formulae in turn and check that they are not already entailed by the set of facts (line 5), but are consistent with respect to it (line 6). If so, we add the new argument to the set of consistent arguments. It terminates by returning the set of consistent arguments. Trivially, MAKE-PSARGS is sound and complete wrt $NMFArgs(\Theta, \Delta)$. Since we enumerate every subset of $\wp(\Theta)$, it is clearly complete, and since we form arguments only from consistent elements of $\wp(\Theta)$, it is sound.

6.4 An Iterative Approach

I earlier suggested that reusing existing software would be a desirable aim. However, when we consider the use of schematic rules with the existing rule engine, there are various technical problems. Specifically, the existing rule engine infers *all* of the literals from a schematic rule, as the following example shows:

Example 6.4.1. Let $\mathcal{K}$ be a vocabulary, $\Delta_\emptyset$ the empty ontology, $\mathcal{R}_\mathcal{K}$ the set of defeasible rules and $\mathcal{G}_\mathcal{K}$ the set of literals. Let $\Theta \subseteq \mathcal{R}_\mathcal{K}$ be a set of rules and $\Gamma \subseteq \mathcal{G}_\mathcal{K}$ be a set of literals. Then $\Omega = (\Theta, \Delta_\emptyset \cup \Gamma)$ be a S-OAF where Θ and Γ are as below:

$$\Theta = \{r_1 : A(x) \Rightarrow B(x), r_2 : B(x) \Rightarrow C(x)\}$$

and

$$\Gamma = \{A(i_1), A(i_2)\}$$

Then we can form the following arguments:

$a_1 \ \langle \{r_1\}, B(i_1) \ \rangle$

$a_2 \ \langle \{r_1\}, B(i_2) \ \rangle$

$a_3 \ \langle \{r_1, r_2\}, C(i_1) \ \rangle$

$a_4 \ \langle \{r_1, r_2\}, C(i_2) \ \rangle$

Note that there are two arguments (a_3 and a_4) with the same support, but with different claims. This is because they represent different groundings of the two schematic rules. So far in the thesis I have assumed that we have some way of distinguishing different groundings of a schematic rule; unfortunately, the available rule-engine does not do this. It returns all the inferences it can construct from a set of rules. Therefore, given $\{r_1\}$, it returns $\{B(i_1), B(i_2)\}$, and given $\{r_1, r_2\}$ it returns $\{B(i_1), B(i_2), C(i_1), C(i_2)\}$. To use this output to make arguments I would then have to do some further processing to divide the output into the claims of arguments. Although possible, I have decided to take a simpler approach. This is based on the observation that each defeasible rule has a single literal in the head. Therefore, the inferences due to any single schematic rule will be a set of literals, each of which is (potentially) the claim of an argument. This also means that, in general, the input to our argument generation algorithm will be a set of literals, an ontology, some rules and some arguments.

There is also an additional possible benefit - since we know that the support of an argument is consistent, if we attempt to extend an argument by adding new rules to the support, the resulting set of rules is more likely to be consistent than if we picked a similarly sized set of rules at random.

Example 6.4.2. Let $\mathcal{K}$ be a vocabulary, $\Delta_\emptyset$ the empty ontology, $\mathcal{R}_\mathcal{K}$ the set of defeasible rules and $\mathcal{G}_\mathcal{K}$ the set of literals. Let $\Theta \subseteq \mathcal{R}_\mathcal{K}$ be a set of rules and $\Gamma \subseteq \mathcal{G}_\mathcal{K}$ be a set of literals. Then $\Omega = (\Theta, \Delta_\emptyset \cup \Gamma)$ be a S-OAF where Θ and Γ are as below:

$$\Theta = \{r_1 : A(x) \Rightarrow B(x), r_2 : B(x) \Rightarrow C(x)\}$$

and

$$\Gamma = \{A(i_1), A(i_2)\}$$

Then on the first iteration, we see that Γ will satisfy r_1, but not r_2:

$$Inf(r_1, \Delta \cup \Gamma) = \{B(i_1), B(i_2)\}$$

We know that each of these could be the claim of an argument, so we can test for consistency, and then form the two arguments:

$a_1 \; \langle \{r_1\}, B(i_1) \, \rangle$

$a_2 \; \langle \{r_1\}, B(i_2) \, \rangle$

We then repeat our cycle, this time with rule r_2, but using the claims of a_1 and a_2 as additional facts:

$Inf(r_2, \Delta_\emptyset \cup \Gamma \cup Claim(a_1)) = \{C(i_1)\}$

$Inf(r_2, \Delta_\emptyset \cup \Gamma \cup Claim(a_2)) = \{C(i_2)\}$

If $\{C(i_1) \cup \Delta_\emptyset \cup \Gamma \cup B(i_1)\}$ and $\{C(i_2) \cup \Delta_\emptyset \cup \Gamma \cup B(i_1)\}$ are consistent, we can form two arguments of the form:

$a_3 \; \langle \{r_1, r_2\}, C(i_1) \, \rangle$

$a_4 \; \langle \{r_1, r_2\}, C(i_2) \, \rangle$

Note that this assumes the same behaviour of the rule engine as we assumed in Inf above; the difference is how we manipulate the input and output. Of course, all of this is only necessary because the available software does not perform exactly as I would like, and so I have to develop an algorithm to handle the problems with the available rule engine. The alternative would be to rewrite the rule-engine entirely, so that it performed as I would like. Unfortunately, given the complexity of the codebase and the time constraints of the thesis, this is not feasible.

6.4.1 Preparatory Considerations

Although the iterative approach is very simple, it requires us to make a few simple modifications to how we think about arguments. Specifically, I need a way of making sure that we are using all the information in an argument when it is used as the basis for a new one, a way of 'tracking' the usage of non-ground arguments, and a 'base case' to start the iterative process. I deal with these in turn below.

Sub-argument Claims

I said above that we needed to ensure that we made use of all the information contained in an argument. Specifically, we cannot only consider the claim of an argument - we also need to consider other literals that are entailed by the subarguments, as this example shows:

Example 6.4.3. Let $\mathcal{K}$ be a vocabulary, $\Delta_\emptyset$ the empty ontology, $\mathcal{R}_\mathcal{K}$ the set of defeasible rules and $\mathcal{G}_\mathcal{K}$ the set of literals. Let $\Theta \subseteq \mathcal{R}_\mathcal{K}$ be a set of rules and $\Gamma \subseteq \mathcal{G}_\mathcal{K}$ be a set of literals. Then $\Omega = (\Theta, \Delta_\emptyset \cup \Gamma)$ be a S-OAF where Θ and Γ are as below:

$\Theta = \{r_1{:}a \Rightarrow b, r_2 \,{:}a{\Rightarrow}c, r_3 \,{:}b \wedge c{\Rightarrow}d\}$

$\Gamma = \{a\}$

Then

a_1 $\langle \{r_1\}, b\rangle$

a_2 $\langle \{r_2\}, c\rangle$

a_3 $\langle \{r_1, r_2, r_3\}, d\rangle$

then clearly a_1, a_2 are sub-arguments of a_3, but we need to consider the claims of both a_1 and a_2 in order to be able to form a_3 .

The simplest way to do this is to generate the 'total claim' of the argument, which is the union of the claims of all the subarguments of some argument and the claim of the argument itself. This is defined below:

Definition 6.4.4. TotalClaim ($\langle A, \phi\rangle$): $\mathcal{A}_\mathcal{K} \mapsto \mathcal{F}_\mathcal{K}$. Let $\langle A_1\phi_1\rangle$, $\langle A_2\phi_2\rangle$.... $\langle A_n\phi_n\rangle$ be the proper subarguments of $\langle A, \phi\rangle$. Then $TotalClaim(\langle A, \phi\rangle)$returns the set of formulae $\{\phi_1, \phi_2, ..., \phi_n, \phi\}$

We can then extend an argument a by considering one rule at a time and using $\Delta \cup TotalClaim(a)$ as a set of facts.

Keeping track of rules

One of the problems with using non-ground rules is that it can be difficult to keep track of which grounding of each rule we have used, as this example shows:

Example 6.4.5. Let $\mathcal{K}$ be a vocabulary, $\Delta_\emptyset$ the empty ontology, $\mathcal{R}_\mathcal{K}$ the set of defeasible rules and $\mathcal{G}_\mathcal{K}$ the set of literals. Let $\Theta \subseteq \mathcal{R}_\mathcal{K}$ be a set of rules and $\Gamma \subseteq \mathcal{G}_\mathcal{K}$ be a set of literals. Then Ω = (Θ, $\Delta_\emptyset \cup \Gamma$) be a S-OAF where Θ and Γ are as below:

Θ = $\{r_1 : R(x, y) \wedge R(y, z) \Rightarrow R(x, z)\}$

Γ = $\{R(i_1, i_2)$, $R(i_2, i_3)$, $R(i_3, i_4)\}$

On the first iteration, we can form an argument:

$\langle \{r_{1a}\}, R(i_1, i_3)\ \rangle$

However, we can see that there is another argument:

$\langle \{r_{1a}, r_{1b}\}, R(i_1, i_4)\ \rangle$

In the example above, the support looks very close to being a repetition, but when developing arguments manually I can index the rules by their grounding and so see that they are different. The available software does not do this. To resolve this I record the rule name and the literal developed from it as a pair, and use this to effectively implement the indexing of ground rules that I have assumed up until this point.

Definition 6.4.6. Let $\mathcal{K}$ be a vocabulary, $\mathcal{R}_\mathcal{K}$ be the set of defeasible rules and $\mathcal{G}_\mathcal{K}$ the set of literals. Also, let $r \in \mathcal{R}_\mathcal{K}$ be a schematic defeasible rule and $\phi \in \mathcal{G}_\mathcal{K}$ be a literal. Then r is of the form $Label$: $Body \Rightarrow Head$, a ***usage*** is pair $(Label, \phi)$ where $Label$ is the label of a (schematic) rule and ϕ is ground literal that is the head of the one of the defeasible rules represented by the schematic rule. The set of all usages is denoted $\mathcal{U}_\mathcal{K}$.

Since we need usage to check whether we have used a grounding of a rule, we need a function to return the usages of an argument.

Definition 6.4.7. Let $\mathcal{K}$ be a vocabulary, $\mathcal{R}_\mathcal{K}$ be the set of defeasible rules and $\mathcal{G}_\mathcal{K}$ the set of literals. Let $\mathcal{U}_\mathcal{K}$ be the set of usages. The function $Usage : \mathcal{A}_\mathcal{K} \mapsto \wp(\mathcal{U}_\mathcal{K})$ returns the set of usages for some argument.

The point of usages is to use them to track the use of groundings of schematic rules; when implementing, we therefore consider sets of usages instead of sets of rule labels in the support.

Example 6.4.8. Let $\mathcal{K}$ be a vocabulary, $\Delta_\emptyset$ the empty ontology, $\mathcal{R}_\mathcal{K}$ the set of defeasible rules and $\mathcal{G}_\mathcal{K}$ the set of literals. Let $\Theta \subseteq \mathcal{R}_\mathcal{K}$ be a set of rules and $\Gamma \subseteq \mathcal{G}_\mathcal{K}$ be a set of literals. Then $\Omega = (\Theta, \Delta_\emptyset \cup \Gamma)$ is a S-OAF where Θ and Γ are as below:

$\Theta = \{r_1 : A(x) \Rightarrow B(x), r_2 : B(x) \Rightarrow C(x)\}$

$\Gamma = \{A(i_1), A(i_2)\}$

Then we can form arguments

$a_1 \ \langle\{r_1\}, B(i_1)\ \rangle$

$a_2 \ \langle\{r_1\}, B(i_2)\ \rangle$

$a_3 \ \langle\{r_1, r_2\}, C(i_1)\ \rangle$

$a_4 \ \langle\{r_1, r_2\}, C(i_2)\ \rangle$

and

$Usage(a_1) = (r_1, B(i_1))$ and $Usage(a_3) = \{(r_1, B(i_1)),(r_2, C(i_2))\}$

Note that for an argument a, where $Usage(a) = \{(L_1,\phi_1), (L_2,\phi_2)...(L_n, \phi_n)\ \}$, $\{L_1, L_2, \ldots ,L_n\} = Support(a)$ and $\{\phi_1, \phi_2, ...,\phi_n\} = TotalClaim(a)$.

Base Argument

An iterative approach must start somewhere, and I start with the ***dummy argument.***

Definition 6.4.9. Let $\mathcal{K}$ be a vocabulary, I be the set of concrete individuals in $\mathcal{K}$, C the set of class names in $\mathcal{K}$ and Δ an ontology. Then the ***dummy argument***, $a_D = \langle\{\},\top(i_1)\ \rangle$ where $\top$ is the tautology predicate $\in C$, $i_1 \in I$ is a concrete individual and $\top(i_1) \in \Delta$. $Support(a_D) = \{\}$, $Usage(A_D) = \{\}$ and $Claim(a_D) = \top(i_1)$.

Note that the claim of the dummy argument adds no new information - for an ontology that contains some individual i_1, $\top(i_1)$ holds. With these details resolved, I present the algorithms for the iterative approach.

6.4.2 Algorithm

My approach uses two algorithms. The first (MAKE-ARGS) incrementally develops arguments, given a set of existing arguments, a set of rules and an ontology. The second ITERATE controls the execution of MAKE-ARGS, accumulates the results and is responsible for terminating the iterative process. I start by presenting MAKE-ARGS:

MAKE-ARGS$(\Delta, \Theta, Args)$

$outArgs \leftarrow \{\}$
for $a \in$ Args
  for $r \in \Theta$
    $Results \leftarrow Inf(\{\text{r}\}, \Delta \cup TotalClaim(a))$
    for $\phi \in Results$
    if $(Label(r), \phi) \notin Usage(a)$
      if $(conflict_{\Delta}(TotalClaim(A), \phi)$ is $false$
        then
          $claim \leftarrow \phi$
          $support \leftarrow Support(a) \cup \{r\}$
          $Usage(\langle\ (support, claim)\ \rangle) \leftarrow \{Usage(a)\} \cup (Label(r), claim)$
          $outArgs \leftarrow \{outArgs\} \cup \langle\ support, claim\ \rangle$
return $outArgs$

The execution of MAKE-ARGS is controlled by ITERATE:

ITERATE$(Ontology, Rules, Args)$

$AllArgs \leftarrow \{\}$
$ArgSet1 \leftarrow \{DummyArg\}$
while $ArgSet1 \neq \emptyset$
  do
    $AllArgs \leftarrow AllArgs \cup ArgSet1$
    $ArgSet1 \leftarrow$ MAKE-ARGS$(Ontology, Rules, ArgSet1)$
return$(AllArgs \setminus DummyArg)$

The major difference is that unlike MAKE-PSARGS, we track the usages as we develop arguments. and $AllArgs$ acts an accumulator to collect the results returned by each iteration of MAKE-ARGS. Although the development of these algorithms was motivated by practical problems in using an existing rule-engine for defeasible reasoning, understanding their computational performance is also important.

6.4.3 Soundness & Completeness

In this section, I shall show that the iterative approach is sound and complete with respect to $NMArgs$. To do so requires some preliminary steps. I shall use the following symbols: let $\mathcal{K}$ be a vocabulary, I the set of concrete individuals in $\mathcal{K}$, Δ an ontology, $\mathcal{R}_{\mathcal{K}}$ the set of defeasible rules and $\mathcal{G}_{\mathcal{K}}$ the set of literals. Let $\Theta \subseteq \mathcal{R}_{\mathcal{K}}$ be a set of rules, $\Gamma \subseteq \mathcal{G}_{\mathcal{K}}$ be a set of literals and $i_1 \in I$ a concrete individual.

Proposition 6.4.10. MAKE-ARGS *is sound wrt to* $NMFArgs$

Proof. I start by recalling the definition of the rule engine function, Inf :

> The ***rule engine*** function $Inf : \wp(\mathcal{R}_{\mathcal{K}}) \times \wp(\mathcal{O}_{\mathcal{K}}) \mapsto \wp(\mathcal{G}_{\mathcal{K}})$ returns the set of all literals ϕ such that:

1. $(e_i, \Delta) \vdash_{Def} \phi$

2. $\Delta \nvdash_{Ont} \phi$

3. There is no literal ϕ_1 such $(e_i, \Delta) \vdash_{Def} \phi_1$ and $(\phi_1 \cup \Delta) \vdash_{Ont} \phi$

and the conditions that hold for each element of $NMFArgs$, $\langle e_i, \phi_i \rangle$:

1. $(A, \Delta) \vdash_{Def} \phi_i$

2. $\Delta \nvdash_{Ont} \phi_i$

3. There is no literal ϕ_j s.t. $(A, \Delta) \vdash \phi_j$ and $\phi_j \cup \Delta \vdash_{Ont} \phi_i$

4. A is not contradictory

Then I need to show that every element of the output of MAKE-ARGS fulfills these criteria. I start by considering an iteration of MAKE-ARGS, and then show that the first iteration a special case of this.

Consider the n^{th} iteration of MAKE-ARGS where the input is Δ, Θ and $Args$ such that $Args \subseteq NMFArgs$. Therefore for any $a \in Args$, $Support(a)$ is a non-contradictory set of rules, and $TotalClaim(a)$ is consistent. We consider each argument $a \in Args$ in turn, and form a set of rules, e, from the support of each argument and each rule $r \in \Theta$ in turn, such that $e = Support(a) \cup \{r\}$. For each e, Inf returns the set of literals $\phi_1 ... \phi_n$ such that $(e, \Delta) \vdash_{Def} \phi_i$, $\Delta \nvdash_{Ont} \phi_i$ and there is no literal ϕ_k such $(e, \Delta) \vdash_{Def} \phi_i$ and $(\phi_i \cup \Delta) \vdash_{Ont} \phi_k$. For each literal in $Inf(e, \Delta)$ such that $Conflict_{\Delta}(TotalClaim(a),\ \phi_i)$ does not hold, e is not contradictory. From the definition of Inf, we know that $\langle e, \phi_i \rangle$ satisfies the first two conditions of $NMFArgs$, and we have shown that e is not contradictory (thus satisfying condition 3). Therefore the output is sound with respect to $NMFArgs$.

If we now consider the first iteration, we see that the input is a set with a single element, the dummy argument a_D. Then $Args = \{a_D\}$ and $TotalClaim(a_D) = \top(i_1)$, but $\top(i_1) \in \Delta$, and so $\Delta \cup \top(i_1) = \Delta$ and hence $\Delta \cup TotalClaim(a) = \Delta$. We consider every $r \in \Theta$, and for each rule generate a set of literals, $Results$. Then for each literal, $\phi \in Results$, we consider whether $Conflict_{\Delta}(TotalClaim(a)$, $\phi)$ holds. $TotalClaim(a_D) = \top(i_1)$, and $\top(i_1) \in \Delta$, and hence $TotalClaim(a) \cup \{\phi\}$ is conflict-free if ϕ is consistent with the ontology. If this is the case, r is not contradictory, and so we can form a new argument, $\langle \{r\}, \phi \rangle$.

Thus the output of the first iteration is a set of arguments, $\langle e_i, \phi_i \rangle$, where for every element, from the definition of Inf, $(e_i, \Delta) \vdash_{Def} \phi_i$, and $\Delta \nvdash_{Ont} \phi_i$, (satisfying conditions 1 and 2 of $NMFArgs$) and e_i is not contradictory (satisfying condition 3). Thus the output of the first iteration is a set of arguments whose members are all elements of $NMFArgs$, and since this is used as the input to the next iteration, the output of subsequent iterations is also sound wrt $NMFArgs$. □

Proposition 6.4.11. ITERATE *is sound wrt* $NMFArgs$

Proof. From the definition above, we need to ensure that the support of every element of the output of ITERATE is also in $NMFArgs$. Since the output of ITERATE is the union of the output of MAKE-ARGS, and MAKE-ARGS is sound, then ITERATE is sound. □

I now move to consider completeness. However, in order to do so, it will be useful to show that MAKE-ARGS has a particular property with respect to the size of the arguments it produces.

Lemma 6.4.12. *If the input to* MAKE-ARGS *contains an argument* a *where* $Size(a) = n$*, then the output of* MAKE-ARGS *contains every argument* a' *s.t.* a *is a proper subargument of* a' *and* $Size(a') = n + 1$.

Proof. This follows very simply from the definition of MAKE-ARGS. We deal with each argument in turn, so we consider an argument, a, where $Size(a) = n$. Then we consider each rule, r, in turn with $TotalClaim(a)$ and the ontology, and for each literal $\phi \in Inf(\{r\}, \Delta \cup TotalClaim(a))$, where ϕ is consistent with $TotalClaim(a)$ and Δ, we form a new argument $b = \langle Support(a) \cup \{r\}, \phi \rangle$. Since $Size(a)$ is given by the number of elements of $Support(a)$, and $Support(b) = Support(a) \cup \{r\}$, $Size(b) = Size(a)$ +1. For every argument a', $Support(a') \subseteq \Theta$. Since we iterate through every element of Θ, for any argument a' such that a is a proper subargument of a', we will eventually generate a'. □

We can extend this lemma to show that if the input to MAKE-ARGS is a set of arguments $a_1 ... a_n$, where for every argument $Size(a_i) = n$, then the output is a set of arguments, where for every argument $a'_1 ... a'_n$ $Size(a'_i) = n + 1$.

Lemma 6.4.13. *If every element of the set of arguments used as input to* MAKE-ARGS *is of size* n*, then every element of the output of* MAKE-ARGS *is of size* $n + 1$

Proof. I have previously shown that the output contains every argument of size $n + 1$. I now show that the output contains only these arguments. I proceed by contradiction.

Consider some rule r and some argument a, where $Size(a) = n$ and a is a member of the set of arguments used as the input to MAKE-ARGS. Consider some argument b where b is in the output of MAKE-ARGS and a is a subargument of b, $b \neq a$ and $Size(a) = Size(b)$. Then $Claim(b) \in Inf(r, \Delta \cup Claim(a))$. However, if $Size(b) = Size(a)$ and a is a subargument of b, $Support(a) = Support(b)$. If $a \neq b$, then $Claim(a) \neq Claim(b)$, else the arguments are identical. However, if $Support(a) = Support(b)$, and $Claim(a) \neq Claim(b)$, then $Claim(a) \cup \Delta \vdash_{Ont} Claim(b)$, or vice versa, as there cannot be any extra defeasible inference. However, the definition of Inf explicitly excludes those literals with a purely strict entailment, and so b would not be returned by MAKE-ARGS. □

I have now shown that if the input to MAKE-ARGS is a set of arguments of size n, the output consists only of the arguments of size $n + 1$, and contains all such arguments of size $n + 1$. I now use these results to prove that ITERATE is complete.

Proposition 6.4.14. ITERATE *is complete wrt* $NMFArgs$

Proof. I start by recalling the definition of completeness wrt $NMFArgs$

> Then Alg ... is ***complete*** with respect to $NMFArgs$ iff for every $a \in NMFArgs(\Theta, \Delta)$, $a \in Alg(\Theta, \Delta)$.

I also recall the definition of NMFArguments

Let Δ be some DL ontology, $\mathcal{G}_{\mathcal{K}}$ the set of all literals and $\mathcal{R}_{\mathcal{K}}$ the defeasible rule language. For some literal $\phi \in \mathcal{G}_{\mathcal{K}}$and a set of rules $A \subseteq \mathcal{R}_{\mathcal{K}}$, then $\langle A, \phi \rangle_N$ is a ***NMFargument*** for ϕ, iff:

1. $(A, \Delta) \vdash_{Def} \phi$
2. $\Delta \nvdash_{Ont} \phi$
3. There is no literal ϕ_j s.t. $(A, \Delta) \vdash \phi_j$ and $\phi_j \cup \Delta \vdash_{Ont} \phi_i$
4. A is not contradictory

From the proofs above, we know that if the input to MAKE-ARGS consists of NMFArguments of size n, then the output consists only of NMFArguments of size n +1, and the output contains every NMFArgument of size $n + 1$. We know that the first iteration of MAKE-ARGS generates every NMFArgument $a_1..a_n$ where $Size(a_i)$ = 1. Therefore the second iteration will generate every NMFArgument s.t. $Size(a)$ = 2, and so on. More generally, for a NMFArgument $a \in NMFArgs$, where $Size(a)$ = n, the output of the n^{th} iteration of MAKE-ARGS will contain a. The arguments given in the construction OAF include a finite set of defeasible rules, Θ. Since every NMFargument has its support drawn from Θ, the support of every argument is a finite set of rules. Therefore, there is some integer, j , such that for all arguments $a_1....a_n \in NMFArgs$, $Size(a_i) < j$. Therefore, after the j^{th} iteration of MAKE-ARGS, every element of $NMFArgs$ will have been present in an iteration of MAKE-ARGS. Since the output of ITERATES is the union of the output of MAKE-ARGS, for a NMFArgument $a \in NMFArgs$, a is in the output of ITERATES. □

Proposition 6.4.15. ITERATE *terminates*

Proof. We have already suggested this in the proof that ITERATE is complete, but refine it here. ITERATE terminates when the output of the i^{th} iteration of MAKE-ARGS matches the output of the $(i + 1)^{th}$ iteration. The maximum size of an argument developed from Θ is the same size as Θ, and there is exactly one argument of this size. Since on the i^{th} iteration the maximum size of the arguments in the output of MAKE-ARGS is i, and the output of each iteration is a complete set of arguments of size i, when i = $|\Theta|$ MAKE-ARGS generates the argument of size $|\Theta|$, denoted a. This is then used as the input to the $(i + 1)^{th}$ iteration. However, there is no argument b s.t. $a \subset b$, and thus the output of the $(i + 1)^{th}$iteration will be the empty set , and thus ITERATE will terminate. □

Corollary 6.4.16. *For every element* $a \in NMFArgs(\Theta, \Delta)$, $Size(a) \leq |\Theta|$

6.4.4 Complexity

In this section, I shall give an upper limit on the time complexity of ITERATE. Since ITERATE repeatedly calls MAKE-ARGS, I start by considering MAKE-ARGS. I shall use the same symbols as above.

MAKE-ARGS is a nested loop with inference and condition checking in the inner loop. Note that we use various functions (such as consistency checking) that may have considerable running time. However, if we assume that the time complexity of these functions is related to the size of the ontology, and that the ontology is much larger than the total claim of an argument then we can assume an approximately constant time for each of these functions. Thus although we may not know the exact time taken for $conflict(\phi_1, \phi_2)$ to return, we can assume that is constant for any two literals in ϕ, and may be approximated by considering the time taken to check the ontology for consistency (which is constant) and hence can be factored out of the analysis.

I therefore reduce the analysis of MAKE-ARGS to counting the number of times that we iterate over inner loop on lines 4-13. Again, we shall consider an OAF with a set of ground rules, Θ, of size n. I start by considering the simplest approach to MAKE-ARGS.

Proposition 6.4.17. *Let Θ be a set of rules, such that $|\Theta| = t$. Let $Args$ be a set of arguments, such that $|Args| = s$. Then if the arguments to* MAKE-ARGS *are Δ, Θ, $Args$, an upper bound on the time complexity of* MAKE-ARGS *is given by $O(t \times s)$.*

Proof. We consider every argument $a \in Args$ in turn and with each, we test at most every rule in Θ. Since $Args$ is of size s, and Θ is of size t, we can iterate no more than $t \times s$ times. Therefore MAKE-ARGS $\approx O(t \times s)$ □

Corollary. *For a particular number of arguments in the input to* MAKE-ARGS*, the worst case time complexity of* MAKE-ARGS *is linear with respect to the size of Θ.*

Given this upper-bound, we can consider the behaviour of ITERATES. I start with an example:

Example 6.4.18. Consider an OAF such that $|\Theta| = 3$. We know that the input to the first iteration is the dummy argument, a_D. Then using the proof above, we know that MAKE-ARGS iterates a maximum of 3 times. The output is therefore a set of at most 3 arguments. Hence the input to next iteration contains three arguments, and so MAKE-ARGS iterates a maximum of 9 times, and the output is at most of size 9, and the at the next iteration MAKE-ARGS iterates at most 27 times.

From the example above we form the following proposition:

Proposition 6.4.19. *For a ruleset of size t, t^i is an upper-bound on the time-complexity of the i^{th} iteration of* MAKE-ARGS

Proof. We know that the time-complexity of MAKE-ARGS for some number of arguments s and number of rules t is in $O(s \times t)$. I proceed by induction.

Base Case: For the first iteration, we know that $s = 1$, and so the output is of size $1 \times t = t^1$.

Iterative Case: If the input to MAKE-ARGS is of size t^{i-1} at the i^{th} iteration, then the output is at most of size $t^{i-1} \times t = t^i$

Taken together, we know from the base case that the output of the first iteration is at most of size t, and that at each iteration the output increases by at most a factor of t. Therefore the output to the i^{th} iteration is at most of size t^i. □

In order to calculate the time complexity of ITERATES, we need only know how many iterations MAKE-ARGS will perform (i.e. a limit on i) and we have an upper limit on the complexity of ITERATES.

However, we already have this result from above, although in a slightly different form. Recall that we know that ITERATE terminates, and I showed that at most, this occurs when $i = t$. Then the maximum number of iterations of MAKE-ARGS is given by t.

Proposition 6.4.20. *For a ruleset of cardinality t, the worst case time complexity of* ITERATES *is in* $O(\frac{t-t^{t+1}}{1-t})$

Proof. From the proof above, the i^{th}step of MAKE-ARGS has an upper bound on its time-complexity of t^i. There are at most t iterations. Therefore the complexity of ITERATE is given by $\sum t^1 + t^2 + t^3 ... t^t = \sum_{i=1}^{t} t^i$, which is the geometric series, whose partial sum is given by $\frac{t(1-t^t)}{1-t}$. □

The main weakness with this estimate is that it assumes that we can construct an argument whose support is every rule in Θ. Since we have assumed in the rest of the thesis that our rules will conflict, this is clearly very unlikely. Nonetheless, the estimate is useful for seeing how the performance of the algorithm varies with respect to the number of rules in Θ, even if we assume (for example) that we may not be likely to generate any argument with more than five rules in the support.

6.5 Software

The algorithms and approaches described above have been used as the basis for a implemented prototype of the system. This makes extensive use of two existing software packages: the Jena RDF API and the Pellet OWL reasoner. The implemented system is designed only to read rules (authored in Jena's existing rule format), construct arguments, check for conflicts between arguments and decide on the preference ordering between arguments based on the the approach described in the first section of this chapter. It does not handle the construction of dialectical trees, and the preference relationships between rules are hard-coded in a data structure. However, it is sufficient to enable the work necessary for the next chapter. The software is written in Java 1.5, and both compiled and source files are available either with this thesis or as a download. For further details on usage and licensing, please see the files that are distributed with the software.

6.6 Summary

In this chapter, I have presented a simple method for describing preferences between rules, and for extending this to sets of rules. I have also introduced an algorithm for argument generation, and given results showing correctness in reference to a simple approach, as well as an upper-bound on the time complexity of this approach. Together with the work in previous chapters, this work forms the basis for a simple prototype implementation that is used to deliver the results in the next chapter.

Chapter 7

Case Study

To further assess OAF, I have developed a case study. The intention is both to demonstrate and assess how far the techniques presented in the thesis go towards achieving my aims. The evaluation of this is done in three ways: the first presents some simple argumentation metrics; the second assesses the correctness of the formal arguments with respect to a medical guideline, and the third discusses the advantages of OAF over existing techniques. I start by describing how the knowledge-base used in the evaluation was generated.

7.1 Defining the Knowledge Base

In chapter 1, I referred to the large number of medical publications, and this has guided the thesis. However, as a reaction to this, many clinicians instead use guidelines, meta-analyses and reviews instead of the primary literature. These secondary sources are so useful that in some cases the primary literature almost disappears from clinical practice. Although this has some advantages, it also weakens the link between clinical trials and clinical practice. This weakening may change the way the evidence is interpreted: for example, if a trial shows the benefit of chemotherapy for women with Stage 1 breast cancer, a guideline may refer to the benefit in early breast cancer (which most would understand as being Stage 1 or Stage 2). The balance between the use of primary and secondary evidence influences how I assess the work. From a clinical point of view it is a question as to whether I intend to produce a system that can in some circumstances produce the 'right' answers, or a system that can reason with the evidence. For the post-surgical treatment of breast cancer, the question of what is 'right' is dominated by a series of meta-analyses published quinquennially in The Lancet, which effectively define current 'standard care', and so by using only these it would be (reasonably) easy to get a system that produced broadly correct answers, which would be supported by most clinicians. However, there are many newer papers that are not included in these meta-analyses, but which clinicians are interested in, even when their conclusions are less strongly supported. Also, the presence of such dominant sources is not guaranteed in other areas of medicine and a system that can represent some of the richness of medical thinking, with its use of conflicting information, is applicable across a wider range of medicine, and helps to bring the primary literature back into clinical practice.

7.1.1 Sources

In thinking about evidence for the case study, our first instinct might be to turn to a database of the primary literature. However, this is problematic - the primary literature is too large, diverse and gives us no easy way to measure the 'correctness' of the results. An alternative is start with a focus and then use the references given in the guideline. To do this, I have used a breast cancer guideline, and selected a limited selection of references (see below for the criteria used). This helps with the scale of the problem and gives me a standard against which I can assess OAF. Medical guidelines are supposed to be based on the evidence they cite, and so I can use the relationship between statements in the guideline and arguments developed from information in the references to assess OAF. The source of the rules is important in determining the relative defeat status between arguments, but there is distinction between a (medical) ***study*** and a (rule's) ***source***. Medical research is (often) organised as distinct clinical trials. Each trial will lead to several publications, each of which I refer to as a ***study***. However, each published paper (study) may contain information on different outcomes and different groups of patients, and I call each of these a ***source***. As described in the last chapter, I concentrate on summarising data on the year of publication, the evidence level, the number of people in the study and the geographical area from where the patients were drawn, and each study may be associated with one or more sources, but each source is derived from only one study, as illustrated below

The Early Breast Cancer Trialist's Collaborative Group published a significant meta-analysis in 2005 [28]. It presented the effects of various treatments on different outcomes for patients with breast cancer. Two effects that were examined were the effect of 2 and 5 years worth of tamoxifen on breast cancer disease-free survival (DFS). However, the information on which these results are based is different, and so the study generates two rules, but each rule is based on a different population described in the study, and so it also generates two sources:

Source: **EBCTCG2005a**
∧ hasDate(EBCTCG2005a,2005)
∧ hasEvidenceLevel(EBCTCG2005a,1a)
∧ hasNumber(EBCTCG2005a,8311)
∧ hasArea(EBCTCG2005a,World)

Source: **EBCTCG2005b**
∧ hasDate(EBCTCG2005b,2005)
∧ hasEvidenceLevel(EBCTCG2005b,1a)
∧ hasNumber(EBCTCG2005b,14336)
∧ hasArea(EBCTCG2005b,World)

The first relates to a five year course of tamoxifen; the second to a two year course. The two sources cover the same area of the world, and have the same evidence level and date, but have different numbers of patients on which they are based; such differences may be important if the rules that they are the source for are involved in arguments that conflict.

7.1.2 Authoring Rules

The National Cancer Institute's (NCI) Breast Cancer Guideline for Health Professionals was used as the medical guideline. This was chosen as it is has been developed by a well-regarded body, is publicly available and has statements that are explicitly linked to the medical literature. In order to make the case study manageable, I have restricted my attention to the section on "Stage I-IIIA and Operable IIIC breast Cancer" and the subsection on "Hormone Therapy" and within this, the paragraphs headed 'Tamoxifen', 'Tamoxifen and Chemotherapy' and 'Tamoxifen Toxic Effects'. Each study referred to in one of these paragraphs was checked to ensure that it dealt with patients with breast cancer, was related to patient outcomes, and provided a clear piece of evidence for an outcome (the last step excluded 3 rather general reviews). The findings of each paper were checked to ensure that they were statistically significant, and related to the assertion made in the guideline. I then used each outcome related to overall survival, breast cancer disease-free survival or the absolute risk or change in risk of a disease or health parameter to write a defeasible rule. Information on the source of each rule came from the study; in some cases the source directly corresponds to the study; in others, the source of a rule comes from a subset of the patients described in the study. The text of the guideline and reasons for exclusions for each reference are given in Appendix A, along with the rules and information on sources. The process of selection is summarised in Fig. 7.1.

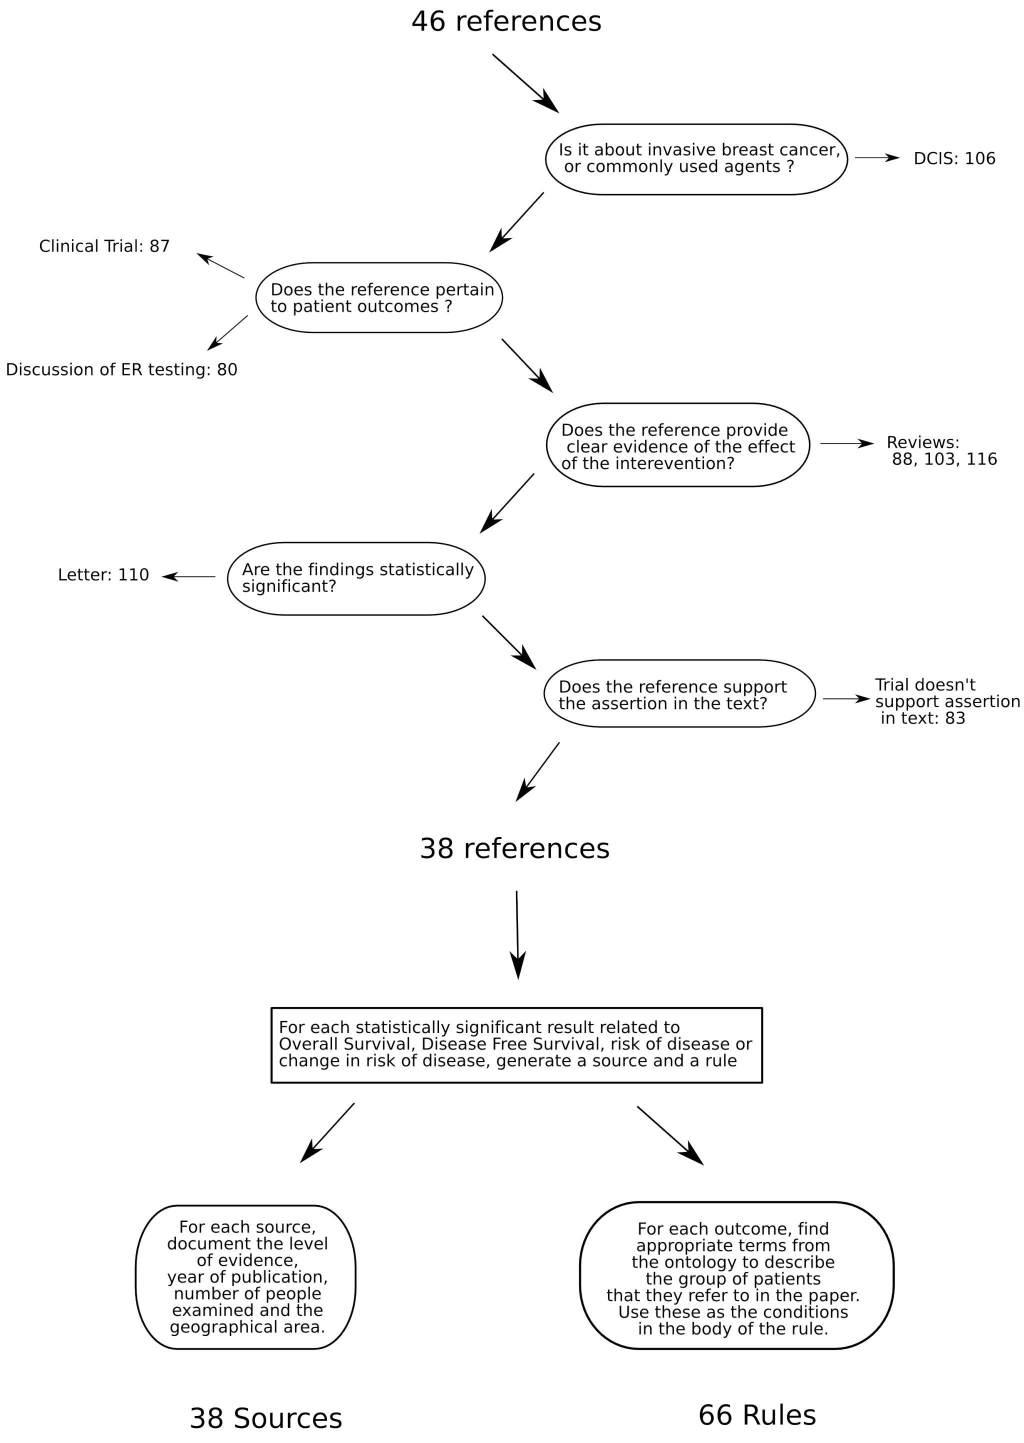

Figure 7.1: Selection of References and Outcomes. 'Acceptable' references are carried downwards; rejected references are filtered to the left and right. Numbers for rejected references are the number of the reference given in the text of the guideline - see Appendix A for details. Acceptable references were then used to generate rules relating to OS, DFS, risk of disease and change in risk of disease. Information about the group of patients on whom these rules were based were used to develop the information in the sources.

7.1.3 KB for evaluation

I then instantiated an R-OAF with the breast cancer ontology, a patient, Ms Jones, a set of values and a set of rules. The complete list of rules and sources is given in Appendix A; facts about Ms Jones and her values are given below.

Example 7.1.1. Then $\Omega_\Delta = (\Theta, \Gamma \cup \Delta_{BC}, V_{MsJones}, R)$ is an R-OAF, where Θ, Φ, Γ, $V_{MsJones}$ and R are as below.

Δ_{BC} is the Breast Cancer Ontology

Θ be the set of rules given in Appendix A and Φ the set of rules such that $\Phi = \{Rewrite(r)$ | $r \in \Theta$ and $Rewrite(\Theta) \neq null\}$

$\mathcal{V}^+_{MsJones}$, $\mathcal{V}^-_{MsJones}$ are sets of values for a patient MsJones and $V = (\mathcal{V}^+_{MsJones}, \mathcal{V}^-_{MsJones})$

$R = [L, Y, N, A]$

Γ = {Women(MsJones), hasDisease(MsJones,ProtoERPosLNPosStage2BrCa), Post-menopausal(MsJones), Aged50Plus(MsJones), hasAge(MsJones, 53)}

$\mathcal{V}^+_{MsJones}$ ={hasDeltaRisk(MsJones, IncreasedOS1.01),
hasDeltaRisk(MsJones, IncreasedOS1.02),
hasDeltaRisk(MsJones, IncreasedOS1.03),
hasDeltaRisk(MsJones, IncreasedOS1.1),
hasDeltaRisk(MsJones, IncreasedOS1.2),
hasDeltaRisk(MsJones, IncreasedBrCaDFS1.05),
hasDeltaRisk(MsJones, IncreasedBrCaDFS1.1),
hasDeltaRisk(MsJones, IncreasedBrCaDFS1.2),
hasDeltaRisk(MsJones, IncreasedBrCaDFS1.21),
hasDeltaRisk(MsJones, IncreasedBrCaDFS1.3),
hasDeltaRisk(MsJones, IncreasedBrCaDFS1.5),
hasDeltaRisk(MsJones, IncreasedBrCaDFS1.7),
hasDeltaRisk(MsJones, DecreasedRiskBrCa0.01),
hasDeltaRisk(MsJones, DecreasedRiskBrCa0.55),
hasDeltaRisk(MsJones, DecreasedRiskBrCa0.6),
hasDeltaRisk(MsJones, DecreasedRiskCardiacDisease0.68),
hasDeltaRisk(MsJones, DecreasedRiskFatalMI0.37),
hasDeltaRisk(MsJones, DecreasedRiskHotFlushes0.8),
}

and

$\mathcal{V}^-_{MsJones}$ ={
hasDeltaRisk(MsJones, DecreasedOS0.7),
hasDeltaRisk(MsJones, DecreasedBrCaDFS0.5),
hasDeltaRisk(MsJones, DecreasedBrCaDFS0.8),

hasDeltaRisk(MsJones, DecreasedEndoCaDFS0.66),

hasDeltaRisk(MsJones, DecreasedEndoCaDSS0.66),

hasDeltaRisk(MsJones, IncreasedRiskCRC1.9),

hasDeltaRisk(MsJones, IncreasedRiskCataract1.14),

hasDeltaRisk(MsJones, IncreasedRiskGastricCa3.0),

hasDeltaRisk(MsJones, IncreasedRiskGastricCa3.2),

hasDeltaRisk(MsJones, IncreasedRiskSideEffects23),

hasDeltaRisk(MsJones, IncreasedRiskPersistentSideEffects2.3),

hasDeltaRisk(MsJones, IncreasedRiskRetinopathy1.52),

hasDeltaRisk(MsJones, IncreasedRiskVenousThrombosis3.4),

hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer2.3),

hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer2.53),

hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer4.1),

hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer6.0),

hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer6.4),

hasDeltaRisk(MsJones, IncreasedRiskEndometrialCancer7.5)

hasDeltaRisk(MsJones, DecreasedOS0.7)}.

Using the prototype described in the last chapter, I generated all the arguments from Ω_{Eval}, and the corresponding dialectical trees for all of them. Preferences between arguments were calculated using the modified method given in the last chapter. This resulted in 730 arguments (and hence 730 argument trees). These run to ~800 pages if printed, and so are available in the electronic supplemental material.

7.2 Some argument metrics

I start with some simple measures of argument and dialectical tree size and structure to provide an overview of the results, although these are rather simplistic, they are also quick and easy to calculate. I then use them to see whether there are any systematic differences between warranted and unwarranted arguments.

7.2.1 Method

The tables below present data on the number of rules in the support of an argument (***Size***), the number of nodes in a dialectical tree (***Tree Size***), the maximum length from a leaf node to the root (***Depth***) and the mean number of children of any non-leaf node (***Av. Branching factor***). For each measure, I present the mean, median, mode and standard deviation, although given the non-normal distribution of the data, the last is for illustrative purposes only.

Results

The case study resulted in the development of 730 formal arguments, which consisted of 151 sets of similar arguments. The tables below present the results for all arguments (Table 7.1), unwarranted (Table 7.2) and warranted arguments (Table 7.3).

	Size	Tree Size	Depth	Av. Branching Factor
Mean	2.65	73.86	1.84	3.67
Median	3	66	2	3.64
Mode	3	139	2	3
SD	0.73	61.8	0.58	1.22

Table 7.1: Summary Statistics for all Arguments. ***Size:*** the number of rules/ argument; ***Tree Size:*** the number of nodes/ tree; ***Depth:*** the number of nodes from root to leaf; ***Av. branching Factor:*** the mean number of attackers per (non-leaf) node.

	Size	Tree Size	Depth	Av. Branching Factor
Mean	2.81	76.98	1.82	3.84
Median	3	66	2	3.89
Mode	3	139	2	4.76
SD	0.55	62.51	0.42	1.04

Table 7.2: Summary Statistics for unwarranted Arguments. ***Size:*** the number of rules/ argument; ***Tree Size:*** the number of nodes/ tree; ***Depth:*** the number of nodes from root to leaf; ***Av. branching Factor:*** the mean number of attackers per (non-leaf) node.

	Size	Tree Size	Depth	Av. Branching Factor
Mean	2.01	60.84	1.9	2.92
Median	2	51	2	3.26
Mode	3	1	2	0
SD	0.99	57.33	1.01	1.58

Table 7.3: Summary Statistics for warranted Arguments. ***Size:*** the number of rules/ argument; ***Tree Size:*** the number of nodes/ tree; ***Depth:*** the number of nodes from root to leaf; ***Av. branching Factor:*** the mean number of attackers per (non-leaf) node.

Discussion

There is a slight trend towards justified arguments having smaller trees and a smaller branching factor. There is a clear difference in the preponderance of trees of size 1 in the set of warranted arguments, reflecting the fact that any argument that has no defeaters will be warranted. Apart from this, however, there is no clear structural difference between trees where the root argument is defeated and those where it is justified. Therefore in general the structure of the tree is not a useful guide to the status of the root node.

Conclusion

This section has presented some simple descriptive statistics about the arguments and dialectical trees. With the exception of trees of size 1, where there is only a root and no leaves, the structure of the tree bears little relation to the justification status of the root node.

7.3 Comparison with clinical guideline

I now consider how we can compare formal and informal arguments. I show that standard techniques for comparison are not applicable, and I therefore briefly introduce a new method. I then use this to compare the informal arguments in the guidelines with the formal arguments developed by OAF.

7.3.1 Introduction

Up until now, I have shown that OAF can produce arguments that seem to match certain intuitive ideas about the way that the results of clinical trials are used. However, I have not shown that they are 'correct'. To do this we need a definition of correctness and since our formal arguments have been developed from references in the guideline, I want to compare the arguments produced by OAF with the guideline. It seems that OAF can fail in two ways - it can either fail to make an argument when we expect one, or we can construct arguments when we do not expect to. It will be helpful to distinguish these two forms of failure and one approach would be to use a 2x2 table (Fig. 7.4).

	Gold Standard Positive	Gold Standard Negative
New Approach Positive	True Positive	False Positive
New Approach Negative	False Negative	True Negative

Table 7.4: Standard 2x2 table

In this work, the informal arguments are the 'gold' standard, and the formal arguments are the new approach. Use of a 2x2 table is based on an assumption that we have a set of instances *all* of which are assessed by both techniques (the 'gold standard' and the new way) and this is not the case here, where we compare formal and informal arguments. Furthermore, a 2x2 table measures error in terms of discordance between the two approaches. If the 'gold standard' is not perfect, some of the 'error' will be a measure of non-erroneous differences. Finally, when we speak of an 'argument' in the guideline, we are using the term informally, or at least in a sense that is not consistent with our formal usage so far - phrases in the guideline are often of the form "For some group of patients, a certain outcome/ treatment is expected/ recommended". An example is given below:

Example 7.3.1. "The overall results of the available evidence suggest that the addition of chemotherapy to tamoxifen in postmenopausal women with ER-positive disease results in a significant, but small, survival advantage"

These informal arguments are clearly different from our formal arguments, and we will therefore need to decide how and when a formal and informal argument can be considered equivalent. A 2x2 table is clearly not a suitable method, and I start by presenting a better approach. The first step is to note that while informal arguments refer to groups of patients, our formal arguments are ground, and I therefore start by choosing a patient (MsJones, who has ER Positive, lymph-node positive breast cancer) as a test case, and generating formal arguments based on the information about her, and then comparing these to the informal arguments applicable to such a patient.

7.3.2 Developing a Method

As well as the different form of the arguments, there is also a problem with the disparity of numbers between the formal (730) and informal (41) arguments. I proceed by dividing both informal and formal arguments according to various criteria, until we are in a position to be clear about how we compare the two. Throughout the following discussion, the key principle is to attempt to capture the expectation given by the informal arguments in formal terms, so we can express whether this expectation is met.

Correspondence of Arguments

The guideline contains informal arguments about the correct treatment option (what we should do) and the effects of various treatments (what we believe will happen if we do it). As described in Chapter 6, our formal arguments are ergonic, commited and hypothetical. Ergonic arguments correspond to informal arguments about the correct treatment options, and hypothetical arguments correspond to the arguments about the effects of different treatments. Committed arguments are different - they are a necessary step in the development of hypothetical arguments. I therefore assess ergonic arguments against the guideline's advice on treatment choice, and hypothetical arguments against our informal arguments about the effect of treatment. Because there is such a close correspondence between ergonic and commited arguments that including both comes close to double-counting an argument, and such types are to some extent an 'artifact' of reasoning in OAF, I shall remove them from the set of formal arguments for my analysis. This approach tells us which types of arguments we should compare. However, it does not tell us whether we should be assessing OAF by looking at its ability to generate arguments or its ability to correctly resolve them, and to answer this we need to study the informal arguments more closely.

What do Informal Arguments mean?

For us to assess the formal arguments against informal ones, we must be clear on the meaning of the informal arguments. It is tempting to see the arguments as being textual representations of formal arguments. This would be incorrect. The informal arguments are more varied than our formal ones, and a little more complex, as this example shows:

Example 7.3.2. The text of an 'informal argument':

> **Text:** 5a: In a trial of node-positive women older than 50 years with hormone-receptor positive tumours, 3-year DFS and OS rates were better in those who received doxorubicin, cyclophosphamide, and tamoxifen versus tamoxifen alone (DFS was 84% vs. 67%; P = .004; OS was 93% vs. 85%; P = .04).[92]
>
> **Conditions:** In women with breast cancer older than 50 years with Hormone-receptor positive tumours

Claim: ACT therapy increases DFS

Even when the claim of the informal argument is simple, it is often a little imprecise. We can capture this by regarding their claim as a class, and checking that the claim of our formal argument is an instance of this class. The informal arguments also have conditions on them, which need to be satisfied for them to hold - In the example above, the information only applies to women over the age of 50 with node-positive, ER-positive disease, but there are other conditions to be met as well. The argument claims that (in this group of women), if they are given ACT, we can expect an increase in breast cancer disease free survival. In assessing the correctness of our formal arguments, we must therefore look for the following:

1. A formal argument that explicitly links ACT and an increase in breast cancer DFS
2. Which is only present if the conditions are satisfied

While an 'incorrect' formal argument would be:

1. One that derived a change in DFS without the application of ACT or
2. One that inferred a different outcome as a result of the use of ACT or
3. One that inferred a reduction in DFS following ACT, but in the wrong group of patients

Therefore, when matching formal and informal arguments, we need to not only consider the claim of the argument, but also the chain of inference by which we generate that claim; I represent this formally, by considering the set of claims of the sub-arguments - as the $TotalClaim()$ function of the last chapter.

Claims vs. Arguments

The second question to address is how we resolve the different numbers of arguments. One obvious approach is to compare individual informal arguments with *groups* of formal ones. However, if we are to do so, we need to ensure that these groups contain similar formal arguments. One approach would be group arguments solely by their claim, but this doesn't meet the criteria I described above, as the example below shows:

Example 7.3.3. Consider two arguments $A = \langle\{r_1, r_2\}, \phi\ \rangle$and $B = \langle\{r_3, r_4\}, \phi\ \rangle$, with sub-arguments A_1= $\langle\{r_1\}, \psi_1\ \rangle$ and B_1 = $\langle\{r_3\}, \psi_2\rangle$. Then we might say that A and B were similar but although they have the same claim, their sub-arguments have different claims.

It should be clear that if we want to analyse arguments that (as above) have a claim of an Increase in DFS, but on the basis that they have a sub-claim of using ACT, grouping arguments by claim alone will not work. I therefore develop a definition of similarity that takes this into account.

Definition 7.3.4. Two arguments A and B are ***similar*** iff $Claim(A)$ = $Claim(B)$ and $TotalClaim(A)$ = $TotalClaim(B)$.

Then we can define what it means for a set of arguments to be similar

Definition 7.3.5. A set of arguments, $Args$, is said to be a set of similar arguments iff every element of $Args$ is similar to every other element.

These definitions provide us with an intuitive method of assembling sets of arguments for assessing against our informal arguments.

Relating Formal & Informal Arguments

We now have a way of producing sets of similar formal arguments to assess against the informal arguments. However, we need to decide whether the members of a set of arguments represent an informal argument. Since the claim of a formal argument is always exactly one ground binary predicate, the first element of which is the patient, we need only specify the class of the second element. We might also want to specify which binary predicate is used, although that seems to be less important. We need to do the same for any sub-claims of the argument, and so for each informal argument we can generate one or more classes for which we expect there to be instances appearing either as the claim of the argument or as the claim of a sub-argument for some argument that is thought to represent the informal argument. The process of constructing the ontology in line with clinical terminology will make this process easier, but it is still not automatic.

Construction vs. Justification

The final consideration is whether we should take the justification status of arguments into consideration when deciding whether informal arguments are justified. Since informal arguments numbered 0 - 9 in table 7.5 conflict with each other, it cannot be the case that each informal argument will be represented by a justified formal argument, and so it seems as though we should assess our informal arguments by considering whether representative formal arguments exist, irrespective of their warranted status. For 10-38, we assume that the patient is taking tamoxifen. However, given this, the arguments are consistent, and so we would expect the resulting arguments to be warranted, and therefore we shall assess the informal arguments against warranted arguments.

7.3.3 Methods

In the section above, I described how we can relate formal and informal arguments, and gave some justification for assessing different arguments in different ways. This is summarised below.

> I start with a set of 730 formal arguments. From this, I remove the committed arguments. We divide the remainder into the set of ergonic arguments, Erg, and epistemic arguments, Epi. In the guideline, we separate the *Treatment Choice* from the other informal arguments, numbered 0 - 38. We proceed as follows:
>
> - An informal argument is ***satisfied*** if its conditions are met by the set of facts relating to MsJones
> - If there is some satisfied informal argument that is represented by a non-empty set of formal arguments, we say that the elements of that set ***correlate*** to the informal argument

- For *Treatment Choice*, we consider the warranted arguments in Erg

- For the informal arguments 0 - 38, we compare them against Epi

- For informal arguments 0- 9, the informal argument is ***represented*** if there exists a non-empty set of similar formal arguments in Epi such that the elements of the set represent the informal argument

- For informal arguments 10 - 38, the argument is ***represented*** if there exists a non-empty set of similar arguments in Epi such that the elements of the set represent the informal argument and are warranted and have a sub-argument with a claim of 2 years of tamoxifen

- Our false positive numbers are given by the number of formal arguments in Epi that are warranted and have a sub-argument with a claim of 2 years of tamoxifen but do not correlate to any informal argument in $0 - 38$

- Our false negative numbers are given by the number of satisifed informal arguments that are not represented

7.3.4 Results

Treatment Choice

I start by considering the ergonic arguments for various treatments. There are 98 of these, consisting of 17 sets of similar arguments. Of these, the only positive ergonic arguments to be justified are those for 2 years worth of tamoxifen. Turning to the conclusion of the guideline, this is apparently correct, although there are some caveats about this which I will discuss later.

Tamoxifen and OS/ DFS

The results are summarised in Table 7.5.

No.	Inf. Claim	Sub-arg. claim	Inf. Sat	F. Sat
0	Tamoxifen increases OS	Tamoxifen	Y	Y
1	No advantage to 10yrs vs. 5 yrs Tamoxifen	> 5 yrs Tamoxifen	Y	Y
2	5 & 10 yrs Tamoxifen are equivalent		Y	N
3	> 5 yrs Tamoxifen does not increase OS	>5 yrs Tamoxifen	N	N
4	Tam. treatment no longer than 5 years		Y‡	Y‡
$5a$	ACT therapy increases DFS	ACT	Y	Y
$5b$	ACT therapy increases OS	ACT	Y	Y
$6a$	CMFT therapy increases DFS	CMFT	N	Y
$6b$	CMFT therapy increases OS		N	Y
7	EarlyCMFT increases DFS	EarlyCMFT	Y	Y
8	CMFT has no effect on DFS	CMFT	Y	Y
9	ChemoTam increases OS	ChemoTam	Y	Y

Table 7.5: Arguments 1 - 9.

Inf. Sat: Whether the conditions on the informal argument are satisfied by Ms. Jones. ***Sub-arg. claim:*** The type of treatment used to achieve the outcome.

F. Sat: Whether there is a set of formal arguments that represent the informal one

Arguments marked with a ‡ are further discussed below.

Informal Arguments: There are 12 informal arguments about the effect of Tamoxifen/ Chemo-Tamoxifen on Overall Survival (OS) and Disease-free Survival (DFS). Of these 12:

- 9 had their conditions met, of which
  - 8 were represented by a set of formal arguments
  - 1 was not
- 3 did not have their conditions met, of which 2 had a formal representation

Formal Arguments: There are 18 sets of similar formal arguments about the effects of Tamoxifen and Chemo-Tamoxifen on OS and DFS. Of these 18:

- 12 have informal correlates
- 6 do not have informal correlates

Side-effects of Tamoxifen

I now consider $10-38$; Recall that all of these assume that the patient is taking 2 years of tamoxifen.

Informal: There are 28 informal arguments about the side-effects of treatment. Of these 28:

- 19 had their conditions met, of which
  - 10 were represented formally and

No.	Inf. Claim	Sub-arg. claim	Inf. Sat	F.Sat
10	Have an increased risk of endometrial cancer		Y	Y
11	Should be evaluated by a gynaecologist◇		N	N
12	Endometrial ca. is higher-	If have Endometrial Ca.	N	N
	grade, more advanced and have a worse outcome			
13	Endometrial cancers are normal cancers	If have Endometrial Ca.	N	N
14	Increased risk of endometrial hyperplasia		Y*	N
15	16% developed hyperplasia		N	N
16	None develop atypical hyperplasia†		N	N
17	Increased risk of GI malignancy		Y	Y
18	Increased risk of DVT		Y	Y
20	Increased risk of endometrial ca.		Y*	Y
21	Reduced Fibrinogen levels		Y	Y
22	Reduced Platelet counts		Y	Y
23	Increased risk of stroke		Y*	N
24	Not an increased risk of stroke		Y*	N
25	10% develop benign ovarian cysts		Y*	N
26	Increased risk of endometrial ca.		Y*	Y
27	Increased risk of gyne. symptoms		Y*	N
28	Clonidine can improve the symptoms	Having hot flushes	N	N
29	Should be carefully assessed◇		N	N
30	Have lower total lipoprotein levels		Y*	N
31	Lower total LDL levels		Y*	N
32	Reduced risk of cardiac disease		Y	Y
33	Reduced risk of cardiac disease (5 yr)	5Yr of Tamoxifen	N	N
34	Reduced risk of fatal MI		Y	Y
35	No Change in CHD Deaths		Y*	N
36	There is a decrease in heart disease		Y	Y
37	Increased Lumbar BMD		Y	N
38	Decreased Lumbar BMD		N	N

Table 7.6: Arguments 10 - 38.

Inf. Sat: Whether the conditions on the informal argument is satisfied by Ms. Jones.

Sub-arg claim: Additional condition or expected claim of sub-argument; 2yrs of Tamoxifen unless otherwise stated.

F. Sat: Whether there is a set of formal arguments that represent the informal one.

◇: Informal Arguments are recommendations about care. †: In patients who do not take Tamoxifen

*: Informal argument differs from published data.

All are discussed below.

  - 9 were not

- 9 did not have their conditions met. None of these were formally represented

Formal: There are 133 sets of formal arguments related to the side-effects of Tamoxifen and Chemo-Tamoxifen. Of these:

- 75 correlated with informal arguments

- 58 did not.

Overall, therefore, 58% of the sets of similar formal have an informal correlate and 64% of the satisfied informal arguments have a formal counterpart.

7.3.5 Discussion

I start by considering the ergonic arguments and note the presence of a rebutted informal argument in the guideline. I then explore the reasons for the errors and I show that there are a variety of causes.

Ergonic arguments

Above, I said that the formal ergonic arguments for 2 years of tamoxifen represented the recommendation in the guideline. However, this is based on a technical, very generous reading of the guideline. In current clinical practice, if a patient such as Ms. Jones were to have tamoxifen, she would receive five years of tamoxifen (not two) and would probably be offered chemotherapy as well. This reason for this difference is instructive. Although there is good evidence to support the use of two years of tamoxifen, there is also evidence that five years of tamoxifen has a greater effect on outcomes (OS and DFS) than two years. Clinical decision making would consider not only the source of the rule, but also the magnitude of the effect - five years worth of tamoxifen has a greater beneficial effect than two years worth, and so is preferred. This difference is not captured in OAF. Positive ergonic arguments are developed as a result of rewriting epistemic rules, based on the fact that they have a positively valued outcome. However, the preference status of the rules is unaffected by rewriting (as they have the same source), and does not consider the magnitude of the outcome. The source of one of the rules about two years worth of tamoxifen is the most preferred source of those considered, and hence arguments based on this are preferred to those for five years of tamoxifen.

This is an interesting result, as my work is not alone in discounting the size of the effect when developing arguments for courses of action. Both Bench-Capon's and Amgoud's work take a similar course. Consideration of effect size is crucial in other approaches to decision-making, such as expected utility approaches. There has been work on this in defeasible logics in the past ([14, 61]) but little in argumentation, and it is interesting to see that the absence of such reasoning has, in this case, resulted in our system producing the incorrect answer. On the other hand, the approach I have followed has delivered a very 'conservative' answer - it has favoured the treatment supported by the strongest evidence, and in this case that is two years worth of tamoxifen.

The other point to consider is that in the set of informal arguments, 4 is a recommendation about a maximum length of tamoxifen treatment. Therefore, there is an argument we should consider it in relation to our ergonic, not epistemic arguments. If we do so, we can see that the warranted ergonic argument for two years worth of tamoxifen is consistent with this. There is a further question about this informal argument, though, which is addressed below.

Informal rebuttal

I mentioned above that the informal argument about the reduction in Anti-Thrombin III is 'rebutted' in the informal guideline by a comment that contradicts the initial claim. This may seem odd, but I suspect that the reason for doing so is related to one of the advantages of argumentation - namely the transparency of the reasoning process. By making, and then undercutting, an argument, the guideline allows the reader to see that we have considered such an argument but then rejected it. This is more transparent than either rejecting the claim outright, or simply not mentioning it, and I think improves the acceptability of the guideline to clinicians.

7.3.6 Reanalysing errors

The analysis above is based on the assumption that that the informal arguments are correct. This allows us to analyse the error rates in the formal arguments, but is open to question. For example, there is a consistent tendency to relax the conditions on the evidence when making inferences about the side-effects of treatments (e.g. $14, 23, 27$). The formal arguments are more adherent to the trial evidence, and therefore in some ways more 'correct' than the informal ones, and this and suggests we should re-examine the assumption that informal arguments are 'correct'.

Treatment & Survival

I start by reconsidering the informal arguments. From above, there are 12 informal arguments about OS and DFS. 3 were not satisfied, of which 2 were still formally represented: $6a$ and $6b$. The informal arguments are based on the NSABP B-20 trial, which studied women with node-negative breast cancer, and since Ms. Jones has node-positive breast cancer, these arguments are not applicable to her. Their formal correlates are based on rules derived from the IBCSG97 trial, which (in part) examined the effect of different CMFT regimes on OS and DFS in post-menopausal women with node-positive disease. This may seem confusing, but this is because the guideline contains various statements, which if taken together could be read as "In women with node-positive or node-negative breast cancer, CMFT increases OS and DFS". Had we made this sort of statement, this new informal argument would have been satisfied, and so the two 'incorrect' sets of formal arguments would no longer be marked as such. However, the methodology outlined above took informal arguments as they were presented in the text, without doing this sort of processing. The other error was for an informal arguments that had its conditions met, but has no formal representation. The reasons for this are different, and are discussed later.

I now consider the formal arguments. From above, there are 18 sets of similar arguments, of which 6 have no informal correlate. I give a single member of each set below, together with the treatment present as part of the sub-claim. In the arguments below, $\mathsf{c} \in \mathcal{R}_{\mathcal{K}}^{\mathcal{C}}$ stands for some commital rule.

a_1 $\langle${IBCSG97a*, c, NCICanada}, hasDeltaRisk(MsJones,ProtoNoChangeOS)$\rangle$ (LateCMFT)

a_2 $\langle${NASBP-B16*, c, NASBP-B16}, hasDeltaRisk(MsJones,IncreasedBrCaDFS1.1)$\rangle$ (ProtoPFTChemoTam)

a_3 ⟨{NASBP-B16*, c, NASBP-B16}, hasDeltaRisk(MsJones,IncreasedBrCaDFS1.3)⟩ (ProtoPAFTChemoTam)

a_4 ⟨{IBCSG97a*, c, NCICanada}, hasDeltaRisk(MsJones,ProtoNoChangeOS)⟩ (EarlyCMFT)

a_5 ⟨ {JNCI1993*, c, EBCTCG2005b}, hasDeltaRisk(MsJones,IncreasedBrCaDFS1.21)⟩ (TamD2Yr)

a_6 ⟨{EBCTCG2005a*, c, EBCTCG2005a}, hasDeltaRisk(MsJones,IncreasedBrCaDFS1.5)⟩ (TamE5Yr)

No.	Inf. Claim	Sub-arg. claim	Inf. Sat	F. Sat
0	Tamoxifen increases OS	Tamoxifen	Y	Y
8	CMFT has no effect on DFS	CMFT	Y	Y
9	ChemoTam increases OS	ChemoTam	Y	Y

Table 7.7: Informal arguments which are approximated by formal arguments

Argument a_1 and a_6 (late CMFT/ early CMFT have no effect on OS) are closely related to 8 but do not represent it. a_2 and a_3 are about the effect of PFT and PAFT on DFS. These are related to 9. a_5 and a_6 are about the effect of Tamoxifen on DFS, and are related to 0. So for every set of formal arguments that has no informal correlate, there is an informal argument which is *nearly* an informal correlate. For four of the six sets, this is because the informal argument mentions OS, while the formal one is about DFS, while for the other two the substitution is the other way around. In both cases, there is an implicit link between a change in DFS and OS. The guideline doesn't make this explicit, as it is a common assumption (in fact violations of this link would be notable), and we can see that this is an example of an implicit inferential step which is neither ontological in nature, nor mentioned in the guideline. One way of describing this would be to treat it as an enthymeme [42], and as such it is a reasonably common one in this domain.

Side-effects of Treatment

Informal Arguments: There are 28 informal arguments about the side-effects of treatment. 19 of these had their conditions met, of which 9 had no formal correlate. In table 7.6 above, note that 10 of the satisfied arguments are marked Y* are those where the informal argument does not strictly match the published data to which it refers. I summarise the differences in the table below:

Inf. Arg No.	Informal Argument	Trial result as published
4	In women with breast cancer	Only applies to LN Neg
	Tam. treatment no longer than 5 years	
14	<NONE>	Applies to women without breast cancer
	Tam. increases the risk of endo. hyperplasia	
20	In women on standard doses of tamoxifen	Not a significant drop in AT III
	Antithrombin III levels are reduced	
23	In women who take tamoxifen	Only for pre-menopausal women
	There is an increased risk of stroke	
24	In women who take tamoxifen	Not statistically significant
	There is not an increased risk of stroke	
25	In women who take tamoxifen	6% in post-menopausal women
	10% develop benign ovarian cysts	
26	In women who take tamoxifen	Node-negative breast cancer
	Increased risk of vasomotor symptoms	
27	In women who take tamoxifen	Node-negative breast cancer
	Increased risk of gynaecologic symptoms	
30	In women who take tamoxifen	Node-negative breast cancer
	They have lower total lipoprotein level	
31	In women who take tamoxifen	Side result of trial
	They have lower low-density lipoprotein level	
35	In women who take Tamoxifen	Not statistically significant
	No change in death from Coronary Heart Disease	

Table 7.8: Problematic Informal Arguments

For each informal argument, the first line refers to to group of people to whom it applies; the second is the conclusion of the argument. The right-hand column summarises the result of the published trial.

Dealing with these in turn, we can start most easily with 23, 25 and 26 and 27. For 23, the referenced evidence only applies to pre-menopausal women, but this restriction is dropped in the informal version. For 26, 27 and 30, the evidence only applies to those with node-negative cancer. For 25, the risk for post-menopausal women should be 6%, not 10%. For 20, the referenced paper states that the reduction in AT III was not of a (clinically) significant level. For 14, the evidence comes from women without breast cancer.

This leaves 24, 31 and 35. In both cases, there was a non-significant change in a non-oncological outcome. Since this was not the main outcome of the trial, it was not used to generate a rule, and there are some methodological reasons for having reservations about using the results of post-hoc analyses such as these. I shall refer to such arguments as being problematic.

Having done this, we can see that in some cases (23, 35) having a deeper information extraction process would probably have solved these problems. Balanced against this, it would have introduced many more arguments not mentioned in the guideline. For 4, 14, 20, 23, 24, 25 and 26 , the guideline makes an assertion that is not explicitly backed by the evidence. The 'relaxation' of the criteria in these cases is clinically plausible, as is the presentation of the risk of ovarian cyst being 10% (as a rounding-up). The mention of AT III levels is a little different, but relates to some other informal arguments that are explicitly weakened soon after they are made, and they will be dealt with in Secn. 7.3.5.

Having been through this process, we can see that for ten of the informal arguments, there are reasons to not expect a formal correlate. It will be helpful to reconsider our analysis including this category of 'problematic' informal arguments, but before we do so, we need to distinguish between cases where we make formal arguments that correspond to informal arguments when the conditions have not been met, and those when we form formal arguments which do not correspond to formal arguments, but which are warranted by the references. We can summarise the results of our reanalysis as follows. We have 28 informal arguments about the side-effects of treatment. Of these 28:

- 19 had their conditions met, of which
  - 9 were unproblematic, of which
    - 8 had formal correlates
    - 1 had no formal correlate
  - 10 were problematic, of which
    - 2 had formal correlates
    - 8 did no formal correlates
- 9 did not have their conditions met. None of these were formally represented

Formal Arguments: I now consider those formal arguments that have no informal correlate. There are 58 such sets of arguments, which I summarise below. Since many of the sets of arguments are closely related, I give a selection of these arguments, the number of sets of similar arguments they represent and what the treatment options range over. In order to summarise the arguments, I use a more compact notation to describe sets of similar argument. For example:

$\langle${JNCI1993, c, Lancet89}, $\phi\rangle$ 18 - TreatmentTypes

shows a single argument, but with the implication that there are 18 sets of similar arguments, all with the same claim, ϕ, as the above argument, and where the claims of the subarguments range over the set of TreatmentTypes.

The formal arguments that have no informal correlate are summarised as follows:

a_1 $\langle${JNCI1993*, c, Lancet89}, hasDeltaRisk(MsJones,DecreasedRiskBrCa0.55)$\rangle$ 6 - Tamoxifen

a_4 $\langle${EBCTCG2005b*, c, JNCITamPrev}, hasDeltaRisk(MsJones,IncreasedRiskCataracts1.14)$\rangle$ 6 - Tamoxifen

a_5 $\langle${EBCTCG1992a*, c, ArchIntMed1991}, hasDeltaRisk(MsJones,IncreasedRiskPersistentSideEffects2.3)$\rangle$ 6 - Tamoxifen

a_6 ⟨{EBCTCG1992a*, c, BMJ1992}, hasDeltaRisk(MsJones,IncreasedRiskRetinopathy1.52)⟩ 6 - Tamoxifen

a_7 ⟨{IBCSG97a*, c, NCICanada}, hasDeltaRisk(MsJones,IncreasedRiskSideEffects23)⟩ 1 - ProtoEarlyCMFTChemoTam

a_8 ⟨{IBCSG97a*, c, NCICanada}, hasDeltaRisk(MsJones,IncreasedRiskSideEffects23)⟩ 1 - ProtoLateCMFTChemTam

a_9 ⟨{IBCSG97a*, c, ECOGVascular}, hasDeltaRisk(MsJones,IncreasedRiskVenousThrombosis3.4)⟩ 1 - ProtoEarlyCMFTChemoTam

a_{10} ⟨{IBCSG97a*, c, ECOGVascular}, hasDeltaRisk(MsJones,IncreasedRiskVenousThrombosis3.4)⟩ 1 - ProtoLateCMFTChemTam

a_{11} ⟨{NASBP-B16*, c, ECOGVascular}, hasDeltaRisk(MsJones,IncreasedRiskVenousThrombosis3.4)⟩ 1 - ProtoPFTChemoTam

a_{12} ⟨{NASBP-B16*, c, ECOGVascular}, hasDeltaRisk(MsJones,IncreasedRiskVenousThrombosis3.4)⟩ 1 - ProtoACT

a_{13} ⟨{NASBP-B16*, c, ECOGVascular}, hasDeltaRisk(MsJones,IncreasedRiskVenousThrombosis3.4)⟩ 1 - ProtoPAFTChemoTam

a_{15} ⟨{JNCI1993*, c, BCRT98}, hasDeltaRisk(MsJones,NoChangeRiskEndometrialAbnormality1)⟩ 6 - Tamoxifen

a_{16} ⟨{Lancet89*, c, Lancet94}, hasDeltaRisk(MsJones,RiskEndometrialAbnormality3.9)⟩ 6 - Tamoxifen

a_{17} ⟨{JNCI1993*, c, RoswellPark}, hasRisk(MsJones,RiskEndometrialCa0.005)⟩ 6 - Tamoxifen

a_{18} ⟨{JNCI1993*, c, GynaeOncol1996}, hasRisk(MsJones,RiskOvarianCyst0.06)⟩ 3 - TamA/TamB/TamC

a_{19} ⟨{EBCTCG2005a*, c, Eye1999}, hasRisk(MsJones,RiskRetinopathy0.12)⟩ 6 - Tamoxifen

For the formal arguments about eye disease, we note that there is an informal argument (29) that mentions the fact "Opthalmic toxic effects have been reported in patients receiving tamoxifen"; our formal arguments report these risks, rather than (as the informal argument does) suggest a course of action. For the risk of persistent-side effects, we see that there is an informal argument (27) talking about similar (though not identical) claims. For the others, however, we are providing new information that the guideline does not contain, and this is hopefully of benefit. In general, these fall into either providing information about the side-effects of Chemo-Tamoxifen (a_7 - a_{13}), providing more detailed information about specific risks (a_{16} - a_{19}) or adding information that does not appear in the guideline (e.g. a_1). In summary, of the formal arguments related to the side-effects of Tamoxifen and Chemo-Tamoxifen, 81 of the 134 had informal correlates, and 53 did not. However, none of the 53 sets of uncorrelated formal arguments were related to unsatisifed informal arguments. Instead, they represent extra information, not found in the guideline.

Unrepresented Arguments: Despite this re-analysis, there are still some arguments which cannot be explained by a closer reading of the evidence. These contain important points in the assessment of OAF and are given below.

> **Text:** 2: A trial that included both node-positive and node-negative women also demonstrated the equivalence of 5 years and 10 years of therapy.

Conditions: In women with breast cancer

Claim: 5 and 10 yrs Tamoxifen are equivalent

Text: 11: Women taking tamoxifen should be evaluated by a gynecologist if they experience any abnormal uterine bleeding

Conditions: Women who take tamoxifen and have abnormal uterine bleeding

Claim: Should be evaluated by a gynaecologist

Text: 29: Ophthalmologic toxic effects have been reported in patients receiving tamoxifen; patients who complain of visual problems should be assessed carefully

Conditions: Women who take tamoxifen and have visual problems

Claim: Should be carefully assessed

There is one informal argument that we fail to represent, 2. However, it is instructive to also consider 11 and 29, which while unsatisfied for this patient also reveal some interesting shortcomings in the formalism.These three informal arguments highlight three separate deficiencies of this work. The second (11) is the easiest to deal with: we could imagine that a course of care would be modelled as an ontological entity, and then we could argue about it as any other. Thus an appointment at a clinic, or a set of appointments, would be another prototypical instance. The third (29) is different. We have many arguments about the risk of certain events, but this type of argument is different. The cursory reading of the informal argument is that women who take tamoxifen may have visual problems. This is incorrect - anyone may have visual problems. But neither is it exactly the same as 'Women who take tamoxifen are at increased risk of visual problems'. Instead, it is saying that in those women who take tamoxifen *and* have visual problems, the tamoxifen may be the *cause* of the visual problems. Opthalmological involvement in oncology is relatively rare, and so an explicit statement linking the two is clinically valuable. In OAF, this sort of argument is very difficult to formalise.

The first is the most difficult, and informative. While our data has come from clinical trials, a cursory glance at the rules shows that they mention only the effect of one agent. This is because I elected to use a simple defeasible rule language as the basis for our argumentation. However, most clinical trials compare the effects of two or more agents against each other, and present their results as 'X had a greater effect than Y'. Encoding such statements in logic is certainly possible, but not trivial, and is of interest not only in our domain but also in many others which would require comparison. The importance for this thesis is that this example neatly highlights a very practical weakness of our formal language.

7.3.7 Conclusions

In this section, I have developed an approach to assessing the performance of my work in relation to an informal guideline, and used this approach to understand where the strengths and weaknesses of both my approach and the guideline lie. I have highlighted two clear shortcomings in the formalism, and shown how some reasonable decisions about the logical framework have lead to deficiencies in reasoning, as well as showing where the guideline authors have made inferences that are not directly sanctioned by the evidence. I have also shown that the extra arguments provided by my system are not erroneous attempts to represent informal arguments in the guideline, but instead represent additional information that is referenced by the guideline but is not quoted by it.

7.4 The advantages of OAF

So far the assessment of OAF has been in terms of its correctness. While some of the work has incidentally used the ontological aspects of OAF, it had not been central to the assessment. However, one of OAF's novelties is that it explicitly uses a DL ontology, and in this section I discuss why I think such an approach is important. The advantages of OAF can be divided into 5 areas:

1. Consistency & Ease
2. Clinical Transparency
3. Domain Suitability
4. Efficiency
5. Separation of Concerns

and I shall deal with them in turn.

Consistency & Ease

For any complex domain there will be multiple authors and maintainers of the rulebase. By defining the language in which our rules are written outside an OAF, we ensure that separately developed rulebases are compatible, both in the syntax and semantics of the predicates in the language. For example, in our case study we use the predicate Breast Cancer(x). Clinically, this could refer to any cancer whose primary site is the breast, or to those cancers that are adenocarcinomas of the breast (we take the second option). Such issues may seem unimportant, but developing a consistent ontological basis for the terms in our domain allows reuse of existing work, as well as allowing the possibility of defining an ontology as a standard in a domain, as SNOMED-CT is becoming within the NHS. Since we allow only the predicates that appear in the ontology to appear in the rules, this consistency is extended to the rules, and it makes the process of authoring rules much easier; it becomes a matter of choosing the appropriate predicates from the ontology. Where we cannot find an appropriate predicate in the ontology, we can define one with reference to existing terms, but then we have a mechanism for checking on the relationships between new terms that authors use.

Clinical Transparency

To ensure that the predicates we use are appropriate for clinicians we use complex class definitions (such as EarlyBreastCancer(x)) to conceal some of the detail from the user - e.g. allowing users to author rules using EarlyBreastCancer(x) instead of enumerating Stage1BreastCancer(x) and Stage2BreastCancer(x).

Thus we construct an ontology which describes our domain, and then define additional predicates (in terms of the original ontology) to capture concepts in the literature, all in DL. Both ontology and rules are understandable to clinicians; the ontology contains the strict inferences that provide our model of the world, while the rules capture our defeasible knowledge, but the use of the ontology allows us to give formal definitions for clinically relevant terms.

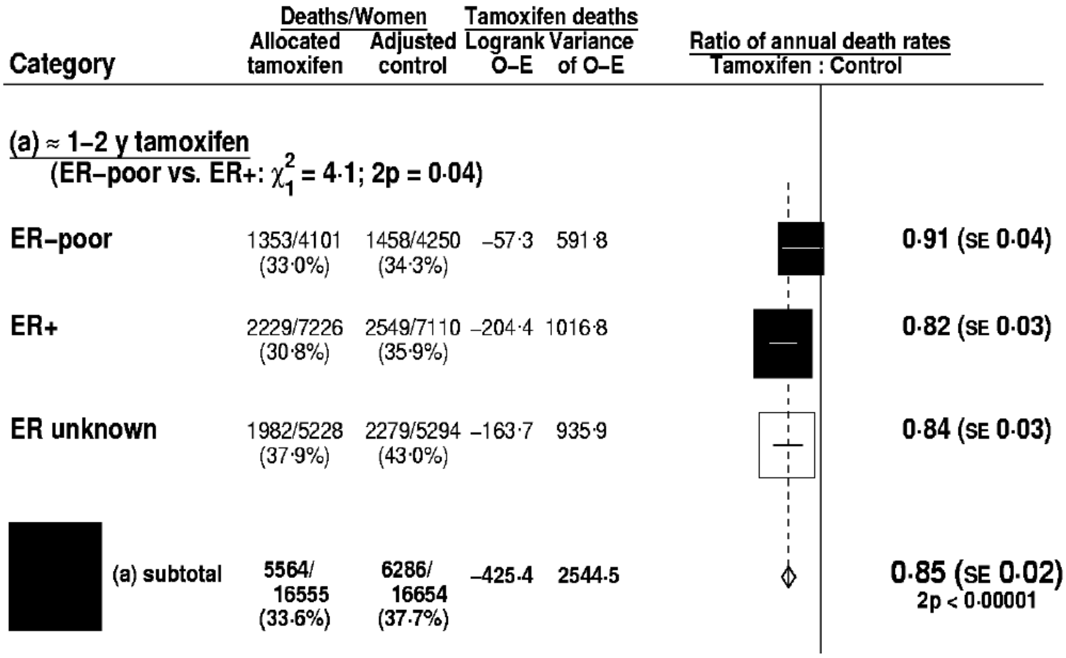

Figure 7.2: Forest Plot for the effect of Tamoxifen

Domain Suitability

Developing the rules may seem difficult, but the way clinical trials are reported in the biomedical literature actually makes this relatively straightforward. Firstly, clinical trials are very explicit about the patients they study (therefore the preconditions of the rules are clearly identifiable); secondly, analyses of the trial results are reported in terms of changes in certain outcomes (therefore the head of the rule is clearly identifiable) and thirdly, meta-analyses of studies have developed methods of summarising many trials in a simple form. For example, the Early Breast Cancer Triallists' Collaborative Group meta-analysis of trials of Tamoxifen included the table in Fig 7.2. Here we see a summary of the results of trials of 1 or 2 years of tamoxifen on survival (measured as annual death rates) in patients with different oestrogen receptor (ER) status. The size of the boxes is proportional to the number of patients studied, and the distance from the central axis represents the size of effect (leftward shift favouring tamoxifen). In this case we can see that tamoxifen has a beneficial effect in those patients whose disease is ER+ve or where ER is unknown, but has no significant effect on patients with ER-ve disease (as the box overlaps the central axis). Diagrams such as these can help one quickly understand the results of the paper. We would argue that clinical trials therefore present their data in a format that is suitable for transformation to a defeasible rule: the patient criteria are the basis of the body of the rule, and the results of the trial form the head. The collection of these terms provide the list of predicates that must appear in the ontology. However, the relations *between* these ontological terms is, in general, not given in the papers - instead, it comes from background clinical knowledge. However, even much of this is codified - for example, the staging of cancers is carefully described in the UICC/AJCC cancer staging manual [2].

Efficiency

Use of the EarlyBreastCancer(x) predicate allows us to subsume those with Stage 1 and Stage 2 disease, while still allowing us to differentiate between them if we need to. In general, if the n terms in a rule each subsume m leaf terms, then the reduction in rules is given by m_1x m_2 x...x m_n . This allows us to relate the structure of the ontology to the expected efficiency gains in the rulebase. In our case study, the efficiency gain is approximately ten-fold from use of ontological predicates, and nearly two-fold from the use of value-based reasoning.

Separation of Concerns

Given that we harness two formalisms, it is tempting to suggest that we simply integrate the two of them into some unifying formalism, and this is what other proposals have done [7]. However, there are both practical and technical reasons for not doing so. Firstly, separation allows the ontology and rulebase to be developed separately, using tools for each as appropriate. Secondly, although reasoning in expressive DLs can be computationally expensive, there are already several optimised reasoners available for DL reasoning. However, argumentative reasoning is even more expensive, and thus any restriction on the number of rules that need to be dealt with by the argumentation engine is desirable. Thirdly, the separation of formalisms allows us to make use of advances in each of the fields as appropriate - for example, the development of more tractable DLs or improved reasoning algorithms [57].

7.5 Conclusion

This chapter has presented a substantial case study, and used it to assess OAF. I have presented some simple metrics on the arguments produced, and developed a methodology to assess the formal arguments against the guideline. I have also presented some more general arguments about why the approach taken in OAF is advantageous.

Of particular interest is the assessment of the formal arguments with respect to the guideline. As part of this, I have shown examples of where the guideline goes beyond that strictly warranted by the evidence it references, demonstrated how the informal arguments use an enthymeme about the relationship between outcomes, and have shown how the assessment of informal vs. formal arguments requires the development of a some new defintions to relate sets of formal arguments to informal ones. The case study also demonstrates a weakness in the ability of the defeasible rule language to capture important aspects of clinical trials, and this may motivate future work in the area.

Chapter 8

Conclusions

This thesis has introduced OAF, a new defeasible-logic based argumentation formalism. The most notable aspect of OAF is its basis in a DL ontology, which defines the predicates that can appear in the rules as well defining whether two literals conflict. I have then provided definitions to extend OAF in order to formalise the practical syllogism and hypothetical reasoning. To assess OAF, I developed a case study based on a breast cancer guideline, and in order to make this feasible I have provided a simple prototype, which necessitated the development of an algorithm to allow for incremental argument generation.

In chapter 1, I said that I aimed to achieve the following:

1. Model the results of clinical trials, and the background knowledge that provides the terms used to describe the results of the trials.
2. Model arguments for both belief and decisions
3. Represent different value-judgements in forming arguments
4. Take a piece of medical knowledge and use it to form different arguments
5. Represent knowledge at different levels of abstraction

and I would suggest that to a greater or lesser extent, I have done so. OAF has two main novel aspects - the incorporation of the ontology and new techniques for developing arguments. In the last chapter, I presented evidence that both elements of the novel work have had some success in achieving their aims, and in addition the comparison of formal and informal arguments highlights some of the problems in guideline development, as well as clearly demonstrating a significant weakness of defeasible rules in representing clinical trial knowledge. OAF is also an interesting comparison to other attempts to formalise the practical syllogism, and it is interesting to note the areas of similarity and difference between the three formalisms discussed in chapter 5.

There are also some less intended consequences of the formalisms in OAF that are worth noting. The most obvious of these is in the definitions of defeasible derivation, where we have the possibility of generating both foreground and background arguments. More interesting than their existence is the fact that, given the clean separation between the ontology and the defeasible rule base, it is hard to prove properties that include background arguments as there are very few restrictions on them. Irrespective of this, I think that the distinction between the two argument types is a useful one, and in many domains, foreground arguments would be preferred for presentation as the background knowledge is considered ‘uninteresting’. The preference scheme introduced in chapter 6 turns out to be far more decisive than that used in DeLP, as it includes the use of the number of patients. This then gives us a very fine-grained approach to deciding on rule preference, which almost inevitably results in one rule (and hence one argument) being clearly preferred to another. In contrast DeLP has techniques for dealing with arguments that are equally preferred, which they they term *blocking defeaters*, which do not occur in the case study. It is an interesting example of how a relatively arbitrary choice (on the basis for deciding on rule preference) can have a profound effect on the behaviour of the system.

The prototype implementation of OAF is clearly too slow to be practical, but is a proof of principle and also illustrates an interesting problem. The upper-bound given in Chapter 6 is too high, but this is largely because it allows for the existence of an argument whose support contains every member of the set of rules, Θ. The rest of the thesis has assumed that the elements of Θ will conflict, and it would be useful to measure the degree of conflict in Θ and use it to produce a parametrised complexity analysis of the algorithm. In particular, different degrees and patterns of conflict (for example, a uniform pattern of conflict which each rule conflicting with a few others, as compared to a few rules conflicting with many others) could lead to the choice of different strategies for reasoning. Although conflict between rules can be difficult to define in a robust sense, as it can be dependent on the set of facts, for a large set of rules, sampling a random subset should lead to a robust estimate of conflict in the entire set without the time involved in processing the entire set.

OAF is deliberately designed to keep the rule-base and ontology separate. In the development of the thesis, I did make an attempt to combine the two, which became very difficult to work with, and in general I think the separation is justified on practical and organisational grounds. However, it does have some negative consequences, and one example of these is the use of functions such as $PeopleNames(\Omega)$, $Treatments(\Omega)$, etc. From an ontological perspective, these are redundant, as it is a quite simple process to determine all the members of any given class, but the separation in OAF means that we have to duplicate them in the argumentation system as well. There are some more subtle problems caused by the mis-match between the expressivity of the logic in the DL ontology and defeasible rule base. Consider the following example:

Example 8.0.1. Let $\mathcal{K}$ be a vocabulary and $\mathcal{G}_{\mathcal{K}}$ the set of literals. Let $\Gamma \subset \mathcal{G}_{\mathcal{K}}$ be a set of literals and let $\Delta_{Wildcat}$ be a DL ontology with axioms as below. $\Omega = (\Theta, \Delta_{Wildcat} \cup \Gamma)$ is a S-OAF where Γ and Θ are as below:

If $\Delta^T_{Wildcat}$ contains the following rules: Leopard(x) $\rightarrow$ Wildcat(x), Lynx(x) $\rightarrow$ Wildcat(x)

and $\Delta^A_{Wildcat}$ = Leopard(Bob) $\vee$ Lynx(Bob)

Then $\Delta_{Wildcat} \vdash_{Ont}$ Wildcat(Bob)

However, if we replace our ontological rules with defeasible rules:

r_1: Leopard(x) $\Rightarrow$ Wildcat(x)

r_2: Lynx(x) $\Rightarrow$ Wildcat(x)

With Θ= $\{r_1, r_2\}$ and Ω= (Θ, Ω), then $\Omega \not\vdash_{def}$ Wildcat(Bob)

The importance of this example is that it shows that the choice as to whether we add a piece of information as a defeasible or DL-Safe rule cannot be an arbitrary one - it makes a difference to the inferences we are then able to make.

Such examples show how the interaction between the formalisms can lead to some unexpected consequences, but also suggest avenues for future work. One option would be to revisit OAF, but using the more expressive DL directly to construct arguments. Given the work by Besnard & Hunter on using first-order logic for argumentation, and recalling the fact that DLs are a subset of first-order logic, one could revise their work to develop a DL-based argumentation framework, taking DL-Safe rules as an obvious starting point. Since DLs are more tractable than first-order logic, this might provide a route to defining a tractable subset of their work. Alternatively, one could replace the defeasible rules in OAF by a defeasible version of DL-Safe rules, and revise OAF. The advantage of this would be that the distinction between the expressivity of the two elements of OAF would lessen.

One of the assumptions early on in the thesis is that the fundamental division into a DL ontology and a set of defeasible rules is a division that can be used in many different domains, and I have not tested this assumption this either in the thesis or outside of it. However, I would suggest that wherever we have a field where there is a complex agreed domain which is unexceptional allied with a defeasible domain where there is some dispute, OAF would seem a reasonable approach.

Other areas of the thesis that seem open to further development include the structure of defeasible rules, and the related work on ergonic and hypothetical reasoning. The structure of the defeasible rules in OAF is taken directly from DeLP, and is common to many other defeasible logic systems. However, it fails to capture important information about the structure of clinical trials, and as I showed in the worked example this has the potential to cause substantial practical problems. The development of a defeasible rule language that was better matched to the domain would therefore be useful for the examples demonstrated, but the problem is not restricted to medicine. Whenever we consider the practical syllogism, we talk about the effects of an action, but there is an assumption about what would otherwise happen if we did nothing. Although there are examples where the alternative is inaction, in many domains we are considering two different courses of action, and a rule language to represent this would seem useful. This richer rule formalism might then allow further work on ergonic and hypothetical reasoning, and one version of a richer rule formalism is currently the subject of a journal paper that has been submitted.

Table of Uniform Notation

The logical language: $\mathcal{K}$

The set of ontological formulae in $\mathcal{K}$: $\mathcal{O}_{\mathcal{K}}$

The set of literals in $\mathcal{K}$: $\mathcal{G}_{\mathcal{K}}$

A set of literals: $\Gamma \subseteq \mathcal{G}_{\mathcal{K}}$

The set of conjunctive formulae in $\mathcal{K}$: $\mathcal{F}_{\mathcal{K}}$

A DL ontology: Δ

The T Box of an ontology Δ^T

The A Box of an ontology Δ^A

The set of all defeasible rules from from elements of $\mathcal{K}$: $\mathcal{R}_{\mathcal{K}}$

A set of rules: $\Theta, \Phi \subseteq \mathcal{R}_{\mathcal{K}}$

For some particular DL ontology, $\Delta \subseteq \mathcal{O}_{\mathcal{K}}$:

C is the set of class names

I is the set of concrete individuals in $\mathcal{K}$, denoted a, b, c, possibly with subscripts

D is the set of datatype value names in $\mathcal{K}$ denoted $d_1 ... d_n$

R is the set of individual-valued role names in $\mathcal{K}$

S is the set of datatype rolenames in $\mathcal{K}$

The sets of rules:

Epistemic rules: $\mathcal{R}_{\mathcal{K}}^{\mathcal{P}}$

Ergonic rules: $\mathcal{R}_{\mathcal{K}}^{\mathcal{P}}$

Commital rules: $\mathcal{R}_{\mathcal{K}}^{\mathcal{P}}$

The sets of arguments:

Epistemic arguments: $\mathcal{A}_{\mathcal{K}}^{\mathcal{P}}$

Ergonic arguments: $\mathcal{A}_{\mathcal{K}}^{\mathcal{R}}$

Commited arguments: $\mathcal{A}_{\mathcal{K}}^{\mathcal{C}}$

Hypothetical arguments: $\mathcal{A}_{\mathcal{K}}^{\mathcal{H}}$

Heuristic arguments: $\mathcal{A}_{\mathcal{K}}^{\mathcal{U}}$

The elements of an OAF:

A set of defeasible rules: Θ

A DL ontology: Δ(see above)

A patient: p

The set of positively-valued patient outcomes for some patient p: $\mathcal{V}_p^+$

The set of negatively-valued patient outcomes for some patient p: $\mathcal{V}_p^-$

The set of values for patient p: $V = \{\mathcal{V}_p^+ \cup \mathcal{V}_p^-\}$

R is a precedence order over rules

The different types of OAF are:

A simple OAF $\Omega = (\Theta,\Delta)$

An extended OAF $\Omega = (\Theta, \Delta, V_p)$

A ranked OAF $\Omega = (\Theta, \Delta, V_p, R)$

Appendix A: Breast Cancer Guideline

The text reproduced below comes from the NCI Breast Cancer PDQ for health professionals, taken in March 2007.

The Guideline

This text was taken from the NCI guideline on breast cancer , accessed on 30th May, 2007. A current version of the guideline is also available, although this differs slightly as it has been updated since the version used below. The guideline is reproduced verbatim apart from some slight formatting changes.

Hormone therapy

If ER status is used to select adjuvant treatment, the study should be performed in a well-established, skilled laboratory. Immunohistochemical assays appear to be at least as reliable as standard ligand-binding assays in predicting response to adjuvant endocrine therapy.[80]

The EBCTCG performed a meta-analysis of systemic treatment of early breast cancer by hormone, cytotoxic, or biologic therapy methods in randomized trials involving 144,939 women with stage I or stage II breast cancer. The most recent analysis, which included information on 80,273 women in 71 trials of adjuvant tamoxifen, was published in 2005.[81] In this analysis, the benefit of tamoxifen was found to be restricted to women with ER-positive or ER-unknown breast tumors. In these women, the 15-year absolute reductions in recurrence and mortality associated with 5 years of use were 12% and 9%, respectively.[81] [Level of evidence: 1iiA] Similar results were found in an International Breast Cancer Study Group trial.[82] Of 1,246 women with stage II disease, only the women with ER-positive disease benefited from tamoxifen.

The most recent EBCTCG meta-analysis also confirmed the benefit of adjuvant tamoxifen in ER-positive premenopausal women.[83] Women younger than 50 years obtained a degree of benefit from 5 years of tamoxifen similar to that obtained by older women. In addition, the proportional reductions in both recurrence and mortality associated with tamoxifen use were similar in women with either node-negative or node-positive breast cancer, but the absolute improvement in survival at 10 years was greater in the latter group (5.3% vs. 12.5% with 5 years of use).[83] [Level of evidence: 1iiA]

The optimal duration of tamoxifen use has been addressed by the EBCTCG meta-analysis and by several large randomized trials. Results from the EBCTCG meta-analysis show a highly significant advantage of 5 years versus 1 to 2 years of tamoxifen with respect to the risk of recurrence (proportionate reduction 15.2%; $P < .001$) and a less significant advantage with respect to mortality (proportionate reduction 7.9%; P = .01). Results from the NSABP B-14 study, which compared 5 years of adjuvant tamoxifen to 10 years of adjuvant tamoxifen for women with early-stage breast cancer, indicate no advantage for continuation of tamoxifen beyond 5 years in women with node-negative, ER-positive breast cancer.[84] [Level of evidence: 1iA] Another trial that included both node-positive and node-negative women also demonstrated the equivalence of 5 years and 10 years of therapy.[85] [Level of evidence: 1iiDi] In both trials, there was a trend toward a worse outcome associated with a longer duration of treatment. In one trial, node-positive women who had already received 5 years of tamoxifen following chemotherapy were randomly assigned to continue therapy or observation.[86] In the ER-positive subgroup, a longer time-to-relapse was associated with continued tamoxifen use, but no improvement in OS was observed.[86] The optimal duration of tamoxifen treatment for node-positive women is still controversial and is being studied in ongoing clinical trials.[87-89] The current recommendation is that adjuvant tamoxifen be discontinued after 5 years in all patients as current standard therapy.[89]

(Refer to the Aromatase inhibitors: letrozole section of this summary for more information on therapy after 5 years of tamoxifen.)

Tamoxifen and chemotherapy

That chemotherapy should add to the effect of tamoxifen in postmenopausal women has been postulated.[90,91] In a trial of node-positive women older than 50 years with hormone receptor - positive tumors, 3-year DFS and OS rates were better in those who received doxorubicin, cyclophosphamide, and tamoxifen versus tamoxifen alone (DFS was 84% vs. 67%; P = .004; OS was 93% vs. 85%; P = .04).[92] [Level of evidence: 1iiA] The NSABP B-20 study compared tamoxifen alone with tamoxifen plus chemotherapy (CMF or sequential methotrexate and 5-FU) in women with node-negative, ER-positive breast cancer. Throughout 5 years of follow-up, the chemotherapy plus tamoxifen regimen resulted in 91% DFS and 96% OS compared with an 87% DFS and 94% OS with tamoxifen alone.[93] [Level of evidence: 1iiA] In another study of postmenopausal women with node-positive disease, tamoxifen alone was compared with tamoxifen plus three different schedules of CMF. A small, DFS advantage was conferred by the addition of early CMF to tamoxifen in women with ER-positive disease.[94] [Level of evidence: 1iiDi] Another study in a similar patient population, in which women were randomized to receive adjuvant tamoxifen with or without CMF, however, showed no benefit in the chemotherapy arm; in this study, intravenous (day 1 every 3 weeks) rather than oral cyclophosphamide was used.[95] [Level of evidence: 1iiA] The overall results of the available evidence suggest that the addition of chemotherapy to tamoxifen in postmenopausal women with ER-positive disease results in a significant, but small, survival advantage.

Tamoxifen toxic effects

The use of adjuvant tamoxifen has been associated with certain toxic effects. The most important is the development of endometrial cancer which, in large clinical trials, has been reported to occur at a rate that is two times to seven times greater than that observed in untreated women.[96-99] Women taking tamoxifen should be evaluated by a gynecologist if they experience any abnormal uterine bleeding. Although one retrospective study raised concern that endometrial cancers in women taking tamoxifen (40 mg/day) had a worse outcome and were characterized by higher-grade lesions and a more advanced stage than endometrial cancers in women not treated with tamoxifen, other larger studies using standard tamoxifen doses (20 mg/day) have not supported this finding.[96,100,101] Similar to estrogen, tamoxifen produces endometrial hyperplasia, which can be a premalignant change. In a cohort of women without a history of breast cancer randomized to receive tamoxifen or placebo on the British Pilot Breast Cancer Prevention Trial, 16% of those on tamoxifen developed atypical hyperplasia at varying times from the start of treatment (range, 3-75 months; median, 24 months), while no cases occurred on the control arm.[102] The value of endometrial biopsy, hysteroscopy, and transvaginal ultrasound as screening tools is unclear.[103,104] Of concern is an increased risk of gastrointestinal malignancy after tamoxifen therapy, but these findings are tentative, and further study is needed.[105]

Tamoxifen is also associated with an increased incidence of deep venous thrombosis and pulmonary emboli. In several adjuvant studies, the incidence ranged from 1% to 2%.[84,92,93,106,107] Clotting factor changes have been observed in controlled studies of prolonged tamoxifen use at standard doses; antithrombin III, fibrinogen, and platelet counts have been reported to be minimally reduced in patients receiving tamoxifen.[108] The relationship of these changes to thromboembolic phenomena is not clear. Tamoxifen may also be associated with an increased risk of strokes.[107,109,110] In the NSABP Breast Cancer Prevention Trial, this increase was not statistically significant.[109]

Another potential problem is the development of benign ovarian cysts, which occurred in about 10% of women in a single study.[111] The relationship between tamoxifen and ovarian tumors requires further study.[112] Short-term toxic effects of tamoxifen may include vasomotor symptoms and gynecologic symptoms (e.g., vaginal discharge or irritation).[113] Clonidine can ameliorate hot flashes in some patients.[114]

Ophthalmologic toxic effects have been reported in patients receiving tamoxifen; patients who complain of visual problems should be assessed carefully.[115-117] Because the teratogenic potential of tamoxifen is unknown, contraception should be discussed with patients who are premenopausal or of childbearing age and are candidates for treatment with this drug.

Tamoxifen therapy may also be associated with certain beneficial estrogenic effects, including decreased total and low-density lipoprotein levels.[118,119] A large controlled Swedish trial has shown a decreased incidence of cardiac disease in postmenopausal women taking tamoxifen. Results were better for women taking tamoxifen for 5 years than for women taking it for 2 years.[120] In another trial, the risk of fatal myocardial infarction was significantly decreased in patients receiving adjuvant tamoxifen for 5 years versus those treated with surgery alone.[119] In the NSABP B-14 study, the annual death rate due to coronary heart disease was lower in the tamoxifen group than in the placebo group (0.62 per 1,000 vs. 0.94 per 1,000), but this difference was not statistically significant.[121] To date, three large controlled trials have shown a decrease in heart disease.[119-121]

Controlled studies have associated long-term tamoxifen use with preservation of bone mineral density of the lumbar spine in postmenopausal women.[122-124] In premenopausal women, decreased bone mineral density is a possibility.[125]

The concluding part of the guideline is a set of treatment suggestions, as laid out below:

Patient Group	Treatment
Premenopausal, ER-positive or PR-positive	Chemotherapy plus tamoxifen
Premenopausal, ER-negative or PR-negative	Chemotherapy
Postmenopausal, ER-positive or PR-positive	Tamoxifen plus chemotherapy, tamoxifen alone
Postmenopausal, ER-negative or PR-negative	Chemotherapy
Older than 70 years	Tamoxifen alone; consider chemotherapy

Table 8.1: Treatment by patient group

The relevant references, with numbers, are given below. This is to the references to be cross-checked with the information in Chapter 7, as the numbering of the references has changed between this version of the guideline and the current one.

References

80: Harvey JM, Clark GM, Osborne CK, et al.: Estrogen receptor status by immunohistochemistry is superior to the ligand-binding assay for predicting response to adjuvant endocrine therapy in breast cancer. J Clin Oncol 17 (5): 1474-81, 1999.

81: Early Breast Cancer Trialists' Collaborative Group (EBCTCG).: Effects of chemotherapy and hormonal therapy for early breast cancer on recurrence and 15-year survival: an overview of the randomised trials. Lancet 365 (9472): 1687-717, 2005.

82: Colleoni M, Gelber S, Goldhirsch A, et al.: Tamoxifen after adjuvant chemotherapy for premenopausal women with lymph node-positive breast cancer: International Breast Cancer Study Group Trial 13-93. J Clin Oncol 24 (9): 1332-41, 2006.

83: Tamoxifen for early breast cancer: an overview of the randomised trials. Early Breast Cancer Trialists' Collaborative Group. Lancet 351 (9114): 1451-67, 1998.

84: Fisher B, Dignam J, Bryant J, et al.: Five versus more than five years of tamoxifen for lymph node-negative breast cancer: updated findings from the National Surgical Adjuvant Breast and Bowel Project B-14 randomized trial. J Natl Cancer Inst 93 (9): 684-90, 2001.

85: Stewart HJ, Prescott RJ, Forrest AP: Scottish adjuvant tamoxifen trial: a randomized study updated to 15 years. J Natl Cancer Inst 93 (6): 456-62, 2001.

86: Tormey DC, Gray R, Falkson HC: Postchemotherapy adjuvant tamoxifen therapy beyond five years in patients with lymph node-positive breast cancer. Eastern Cooperative Oncology Group. J Natl Cancer Inst 88 (24): 1828-33, 1996.

87: Baum M, Cancer Research UK Clinical Groups at Guy's King's & St. Thomas' Hospitals : Phase III Randomized Study of Adjuvant Therapy with 2 vs 5 Years of Tamoxifen, with Optional Randomization for Radiotherapy, in Breast Cancer Patients Over the Age of 50 (Summary Last Modified 03/98), CRC-PHASE-III-86003, Clinical trial, Closed.

88: Bilimoria MM, Assikis VJ, Jordan VC: Should Adjuvant Tamoxifen Therapy be Stopped at 5 Years? Cancer J Sci Am 2 (3): 140-50, 1996.

89: Swain SM: Tamoxifen: the long and short of it. J Natl Cancer Inst 88 (21): 1510-2, 1996.

90: Systemic treatment of early breast cancer by hormonal, cytotoxic, or immune therapy. 133 randomised trials involving 31,000 recurrences and 24,000 deaths among 75,000 women. Early Breast Cancer Trialists' Collaborative Group. Lancet 339 (8784): 1-15, 1992.

91: Systemic treatment of early breast cancer by hormonal, cytotoxic, or immune therapy. 133 randomised trials involving 31,000 recurrences and 24,000 deaths among 75,000 women. Early Breast Cancer Trialists' Collaborative Group. Lancet 339 (8785): 71-85, 1992.

92: Fisher B, Redmond C, Legault-Poisson S, et al.: Postoperative chemotherapy and tamoxifen compared with tamoxifen alone in the treatment of positive-node breast cancer patients aged 50 years and older with tumors responsive to tamoxifen: results from the National Surgical Adjuvant Breast and Bowel Project B-16. J Clin Oncol 8 (6): 1005-18, 1990.

93: Fisher B, Dignam J, Wolmark N, et al.: Tamoxifen and chemotherapy for lymph node-negative, estrogen receptor-positive breast cancer. J Natl Cancer Inst 89 (22): 1673-82, 1997.

94: Effectiveness of adjuvant chemotherapy in combination with tamoxifen for node-positive postmenopausal breast cancer patients. International Breast Cancer Study Group. J Clin Oncol 15 (4): 1385-94, 1997.

95: Pritchard KI, Paterson AH, Fine S, et al.: Randomized trial of cyclophosphamide, methotrexate, and fluorouracil chemotherapy added to tamoxifen as adjuvant therapy in postmenopausal women with node-positive estrogen and/or progesterone receptor-positive breast cancer: a report of the National Cancer Institute of Canada Clinical Trials Group. Breast Cancer Site Group. J Clin Oncol 15 (6): 2302-11, 1997.

96: Fornander T, Rutqvist LE, Cedermark B, et al.: Adjuvant tamoxifen in early breast cancer: occurrence of new primary cancers. Lancet 1 (8630): 117-20, 1989.

97: Magriples U, Naftolin F, Schwartz PE, et al.: High-grade endometrial carcinoma in tamoxifen-treated breast cancer patients. J Clin Oncol 11 (3): 485-90, 1993.

98: Fisher B, Costantino JP, Redmond CK, et al.: Endometrial cancer in tamoxifen-treated breast cancer patients: findings from the National Surgical Adjuvant Breast and Bowel Project (NSABP) B-14. J Natl Cancer Inst 86 (7): 527-37, 1994.

99: van Leeuwen FE, Benraadt J, Coebergh JW, et al.: Risk of endometrial cancer after tamoxifen treatment of breast cancer. Lancet 343 (8895): 448-52, 1994.

100: Barakat RR, Wong G, Curtin JP, et al.: Tamoxifen use in breast cancer patients who subsequently develop corpus cancer is not associated with a higher incidence of adverse histologic features. Gynecol Oncol 55 (2): 164-8, 1994.

101: Cuenca RE, Giachino J, Arredondo MA, et al.: Endometrial carcinoma associated with breast carcinoma: low incidence with tamoxifen use. Cancer 77 (10): 2058-63, 1996.

102: Kedar RP, Bourne TH, Powles TJ, et al.: Effects of tamoxifen on uterus and ovaries of postmenopausal women in a randomised breast cancer prevention trial. Lancet 343 (8909): 1318-21, 1994.

103: Neven P, De Muylder X, Van Belle Y, et al.: Tamoxifen and the uterus. Br Med J 309 (6965): 1313-1314, 1994.

104: Bertelli G, Venturini M, Del Mastro L, et al.: Tamoxifen and the endometrium: findings of pelvic ultrasound examination and endometrial biopsy in asymptomatic breast cancer patients. Breast Cancer Res Treat 47 (1): 41-6, 1998.

105: Rutqvist LE, Johansson H, Signomklao T, et al.: Adjuvant tamoxifen therapy for early stage breast cancer and second primary malignancies. Stockholm Breast Cancer Study Group. J Natl Cancer Inst 87 (9): 645-51, 1995.

106: Fisher B, Dignam J, Wolmark N, et al.: Tamoxifen in treatment of intraductal breast cancer: National Surgical Adjuvant Breast and Bowel Project B-24 randomised controlled trial. Lancet 353 (9169): 1993-2000, 1999.

107: Saphner T, Tormey DC, Gray R: Venous and arterial thrombosis in patients who received adjuvant therapy for breast cancer. J Clin Oncol 9 (2): 286-94, 1991.

108 Love RR, Surawicz TS, Williams EC: Antithrombin III level, fibrinogen level, and platelet count changes with adjuvant tamoxifen therapy. Arch Intern Med 152 (2): 317-20, 1992.

109: Fisher B, Costantino JP, Wickerham DL, et al.: Tamoxifen for prevention of breast cancer: report of the National Surgical Adjuvant Breast and Bowel Project P-1 Study. J Natl Cancer Inst 90 (18): 1371-88, 1998.

110: Dignam JJ, Fisher B: Occurrence of stroke with tamoxifen in NSABP B-24. Lancet 355 (9206): 848-9, 2000.

111: Shushan A, Peretz T, Uziely B, et al.: Ovarian cysts in premenopausal and postmenopausal tamoxifen-treated women with breast cancer. Am J Obstet Gynecol 174 (1 Pt 1): 141-4, 1996.

112: Cohen I, Beyth Y, Tepper R, et al.: Ovarian tumors in postmenopausal breast cancer patients treated with tamoxifen. Gynecol Oncol 60 (1): 54-8, 1996.

113: Love RR, Cameron L, Connell BL, et al.: Symptoms associated with tamoxifen treatment in postmenopausal women. Arch Intern Med 151 (9): 1842-7, 1991.

114: Goldberg RM, Loprinzi CL, O'Fallon JR, et al.: Transdermal clonidine for ameliorating tamoxifen-induced hot flashes. J Clin Oncol 12 (1): 155-8, 1994.

115: Bentley CR, Davies G, Aclimandos WA: Tamoxifen retinopathy: a rare but serious complication. BMJ 304 (6825): 495-6, 1992.

116: Nayfield SG, Gorin MB: Tamoxifen-associated eye disease. A review. J Clin Oncol 14 (3): 1018-26, 1996.

117: Noureddin BN, Seoud M, Bashshur Z, et al.: Ocular toxicity in low-dose tamoxifen: a prospective study. Eye 13 (Pt 6): 729-33, 1999.

118: Love RR, Wiebe DA, Newcomb PA, et al.: Effects of tamoxifen on cardiovascular risk factors in postmenopausal women. Ann Intern Med 115 (11): 860-4, 1991.

119: McDonald CC, Stewart HJ: Fatal myocardial infarction in the Scottish adjuvant tamoxifen trial. The Scottish Breast Cancer Committee. BMJ 303 (6800): 435-7, 1991.

120: Rutqvist LE, Mattsson A: Cardiac and thromboembolic morbidity among postmenopausal women with early-stage breast cancer in a randomized trial of adjuvant tamoxifen. The Stockholm Breast Cancer Study Group. J Natl Cancer Inst 85 (17): 1398-406, 1993.

120: Costantino JP, Kuller LH, Ives DG, et al.: Coronary heart disease mortality and adjuvant tamoxifen therapy. J Natl Cancer Inst 89 (11): 776-82, 1997.

121: Love RR, Mazess RB, Barden HS, et al.: Effects of tamoxifen on bone mineral density in postmenopausal women with breast cancer. N Engl J Med 326 (13): 852-6, 1992.

122: Kristensen B, Ejlertsen B, Dalgaard P, et al.: Tamoxifen and bone metabolism in postmenopausal low-risk breast cancer patients: a randomized study. J Clin Oncol 12 (5): 992-7, 1994.

123: Love RR, Barden HS, Mazess RB, et al.: Effect of tamoxifen on lumbar spine bone mineral density in postmenopausal women after 5 years. Arch Intern Med 154 (22): 2585-8, 1994.

124: Powles TJ, Hickish T, Kanis JA, et al.: Effect of tamoxifen on bone mineral density measured by dual-energy x-ray absorptiometry in healthy premenopausal and postmenopausal women. J Clin Oncol 14 (1): 78-84, 1996.

Appendix B: Rules

The following rules were generated from the studies referenced in the guideline.

EBCTCG2005a: Women(p) ∧ hasDisease(p,q) ∧ EarlyBreastCancer(q) ∧ ERPos(q)
∧ hasTreatment(p,s) ∧ Tamoxifen5YrCourse(s)
∧ IncreasedBrCaDFS(t) ∧ hasValue(t,1.5)
⇒ hasDeltaRisk(p,t)

EBCTCG2005b: Women(p) ∧ hasDisease(p,q) ∧ EarlyBreastCancer(q) ∧ ERPos(q)
∧ hasTreatment(p,s) ∧ Tamoxifen2YrCourse(s)
∧ IncreasedBrCaDFS(t) ∧ hasValue(t,1.21)
⇒ hasDeltaRisk(p,t)

IBCSGT13-93a: Women(p) ∧ hasDisease(p,q) ∧ BreastCancer(q) ∧ LNPos(q) ∧ ERPos(q)
∧ Premenopausal(p)
∧ hasTreatment(p,s) ∧ Tamoxifen5YrCourse(s)
∧ IncreasedBrCaDFS(t) ∧ hasValue(t,1.7) ∧ refersDisease(t,q)
⇒ hasDeltaRisk(p,t)

IBCSGT13-93a: Women(p) ∧ hasDisease(p,q) ∧ BreastCancer(q) ∧ LNPos(q) ∧ ERNeg(q)
∧ Premenopausal(p)
∧ hasTreatment(p,s) ∧ Tamoxifen5YrCourse(s)
∧ DecreasedBrCaDFS(t) ∧ hasValue(t,0.5) ∧ refersDisease(t,q)
⇒ hasDeltaRisk(p,t)

IBCSGT13-93a: Women(p) ∧ hasDisease(p,q) ∧ BreastCancer(q) ∧ LNPos(q) ∧ ERUnknown(q) ∧ Premenopausal(p)
∧ hasTreatment(p,s) ∧ Tamoxifen5YrCourse(s)
∧ NoChangeBrCaDFS(t) ∧ hasValue(t,0) ∧ refersDisease(t,q)
⇒ hasDeltaRisk(p,t)

IBCSGT13-93a: Women(p) ∧ hasDisease(p,q) ∧ BreastCancer(q) ∧ LNPos(q) ∧ ERPos(q)

∧ Premenopausal(p) ∧ Amenorrhic(p)

∧ hasTreatment(p,s) ∧ Tamoxifen5YrCourse(s)

∧ IncreasedBrCaDFS(t) ∧ hasValue(t,1.7) ∧ refersDisease(t,q)

⇒ hasDeltaRisk(p,t)

NSABP14a: Women(p) ∧ hasDisease(p,q) ∧ BreastCancer(q) ∧ LNNeg(q)

∧ hasTreatment(p,s) ∧ TamoxifenMoreThan5YrCourse(s)

∧ DecreasedBrCaDFS(t) ∧ hasValue(t,0.8) ∧ refersDisease(t,q)

⇒ hasDeltaRisk(p,t)

ScotTamTrial: Women(p) ∧ hasDisease(p,q) ∧ BreastCancer(q)

∧ hasTreatment(p,s) ∧ TamoxifenMoreThan5YrCourse(s)

∧ NoChangeBrCaDFS(t) ∧ hasValue(t,0) ∧ refersDisease(t,q)

⇒ hasDeltaRisk(p,t)

ScotTamTrial: Women(p) ∧ hasDisease(p,q) ∧ BreastCancer(q)

∧ hasTreatment(p,s) ∧ TamoxifenMoreThan5YrCourse(s)

∧ NoChangeOS(t) ∧ hasValue(t,0)

⇒ hasDeltaRisk(p,t)

ECOGa: Women(p) ∧ hasDisease(p,q) ∧ BreastCancer(q) ∧ LNPos(q)

∧ hasTreatment(p,s) ∧ TamoxifenMoreThan5YrCourse(s)

∧ NoChangeOS(t) ∧ hasValue(t,0)

⇒ hasDeltaRisk(p,t)

ECOGa: Women(p) ∧ hasDisease(p,q) ∧ BreastCancer(q) ∧ LNPos(q)

∧ hasTreatment(p,s) ∧ TamoxifenMoreThan5YrCourse(s)

∧ NoChangeBrCaDFS(t) ∧ hasValue(t,0) ∧ refersDisease(t,q)

⇒ hasDeltaRisk(p,t)

JNCIEd: Women(p) ∧ hasDisease(p,q) ∧ BreastCancer(q) ∧ LNNeg(q)

∧ ∧ TamoxifenMoreThan5YrCourse(s)

∧

⇒ ¬ hasTreatment(p,s)

NASBP-B16: Women(p) $\wedge$ hasDisease(p,q) $\wedge$ BreastCancer(q) $\wedge$ Age50Plus(p)
$\wedge$ hasTreatment(p,s) $\wedge$ ACTChemoTam(s)
$\wedge$ IncreasedBrCaDFS(t) $\wedge$ hasValue(1.2) $\wedge$ refersDisease(t,q)
$\Rightarrow$ hasDeltaRisk(p,t)

NASBP-B16: Women(p) $\wedge$ hasDisease(p,q) $\wedge$ BreastCancer(q) $\wedge$ Age50Plus(p)
$\wedge$ hasTreatment(p,s) $\wedge$ ACTChemoTam(s)
$\wedge$ IncreasedOS(t) $\wedge$ hasValue(1.1)
$\Rightarrow$ hasDeltaRisk(p,t)

NASBP-B16: Women(p) $\wedge$ hasDisease(p,q) $\wedge$ BreastCancer(q) $\wedge$ Age50Plus(p)
$\wedge$ hasTreatment(p,s) $\wedge$ PAFTChemoTam(s)
$\wedge$ IncreasedBrCaDFS(t) $\wedge$ hasValue(1.3) $\wedge$ refersDisease(t,q)
$\Rightarrow$ hasDeltaRisk(p,t)

NASBP-B16: Women(p) $\wedge$ hasDisease(p,q) $\wedge$ BreastCancer(q) $\wedge$ Age50Plus(p)
$\wedge$ hasTreatment(p,s) $\wedge$ PFTChemoTam(s)
$\wedge$ IncreasedBrCaDFS(t) $\wedge$ hasValue(1.1) $\wedge$ refersDisease(t,q)
$\Rightarrow$ hasDeltaRisk(p,t)

NSABP20a: Women(p) $\wedge$ hasDisease(p,q) $\wedge$ BreastCancer(q) $\wedge$ ERPos(q) $\wedge$ LNNeg(q)
$\wedge$ hasTreatment(p,s) $\wedge$ MFTChemoTam(s)
$\wedge$ IncreasedBrCaDFS(t) $\wedge$ hasValue(t,1.05) $\wedge$ refersDisease(t,q)
$\Rightarrow$ hasDeltaRisk(p,t)

NSABP20a: Women(p) $\wedge$ hasDisease(p,q) $\wedge$ BreastCancer(q) $\wedge$ ERPos(q) $\wedge$ LNNeg(q)
$\wedge$ hasTreatment(p,s) $\wedge$ CMFTChemoTam(s)
$\wedge$ IncreasedBrCaDFS(t) $\wedge$ hasValue(t,1.05) $\wedge$ refersDisease(t,q)
$\Rightarrow$ hasDeltaRisk(p,t)

NSABP20a: Women(p) $\wedge$ hasDisease(p,q) $\wedge$ BreastCancer(q) $\wedge$ ERPos(q) $\wedge$ LNNeg(q)
$\wedge$ hasTreatment(p,s) $\wedge$ MFTChemoTam(s)
$\wedge$ IncreasedOS(t) $\wedge$ hasValue(t,1.03)
$\Rightarrow$ hasDeltaRisk(p,t)

NSABP20a: Women(p) $\wedge$ hasDisease(p,q) $\wedge$ BreastCancer(q) $\wedge$ ERPos(q) $\wedge$ LNNeg(q)
$\wedge$ hasTreatment(p,s) $\wedge$ CMFTChemoTam(s)

$\wedge$ IncreasedOS(t) $\wedge$ hasValue(t,1.02)

$\Rightarrow$ hasDeltaRisk(p,t)

ICBSG97a: Women(p) $\wedge$ hasDisease(p,q) $\wedge$ BreastCancer(q) $\wedge$ Postmenopausal(p) $\wedge$ LNPos(q)

$\wedge$ hasTreatment(p,s) $\wedge$ EarlyCMFTChemoTam(s)

$\wedge$ IncreasedBrCaDFS(t) $\wedge$ hasValue(t,1.1) $\wedge$ refersDisease(t,q)

$\Rightarrow$ hasDeltaRisk(p,t)

ICBSG97a: Women(p) $\wedge$ hasDisease(p,q) $\wedge$ BreastCancer(q) $\wedge$ Postmenopausal(p) $\wedge$ LNPos(q)

$\wedge$ hasTreatment(p,s) $\wedge$ LateCMFTChemoTam(s)

$\wedge$ NoChangeBrCaDFS(t) $\wedge$ hasValue(t,1.0) $\wedge$ refersDisease(t,q)

$\Rightarrow$ hasDeltaRisk(p,t)

ICBSG97a: Women(p) $\wedge$ hasDisease(p,q) $\wedge$ BreastCancer(q) $\wedge$ Postmenopausal(p) $\wedge$ LNPos(q) $\wedge$ ERPos(q)

$\wedge$ hasTreatment(p,s) $\wedge$ CMFTChemoTam(s)

$\wedge$ IncreasedBrCaDFS(t) $\wedge$ hasValue(t, 1.1) $\wedge$ refersDisease(t,q)

$\Rightarrow$ hasDeltaRisk(p,t)

NCICanada: Women(p) $\wedge$ hasDisease(p,q) $\wedge$ BreastCancer(q) $\wedge$ Postmenopausal(p) $\wedge$ LNPos(q) $\wedge$ HRPos(q)

$\wedge$ hasTreatment(p,s) $\wedge$ CMFTChemoTam(s)

$\wedge$ IncreasedSideEffects(t) $\wedge$ hasValue(t,23)

$\Rightarrow$ hasDeltaRisk(p,t)

NCICanada: Women(p) $\wedge$ hasDisease(p,q) $\wedge$ BreastCancer(q) $\wedge$ Postmenopausal(p) $\wedge$ LNPos(q) $\wedge$ HRPos(q)

$\wedge$ hasTreatment(p,s) $\wedge$ CMFTChemoTam(s)

$\wedge$ OverallSurvival(t) $\wedge$ hasValue(t,0.0)

$\Rightarrow$ hasDeltaRisk(p,t)

Lancet89: Women(p) $\wedge$ hasDisease(p,q) $\wedge$ BreastCancer(q) $\wedge$ Postmenopausal(p)

$\wedge$ hasTreatment(p,s) $\wedge$ Tamoxifen(s)

$\wedge$ ReducedRiskBrCa(t) $\wedge$ hasValue(t, 0.55) $\wedge$ refersDisease(t,q)

⇒ hasDeltaRisk(p,t)

Lancet89: Women(p) ∧ hasDisease(p,q) ∧ BreastCancer(q) ∧ Postmenopausal(p)

∧ hasTreatment(p,s) ∧ Tamoxifen(s)

∧ IncreasedRiskEndoCa(t) ∧ hasValue(t, 6.4) ∧ refersDisease(t,q)

⇒ hasDeltaRisk(p,t)

JCOHighGradeCaTam: Women(p) ∧ hasDisease(p,q) ∧ BreastCancer(q) ∧ hasDisease(p,r)

∧ EndometrialCa(r)

∧ hasTreatment(p,s) ∧ Tamoxifen(s)

∧ PoorPrognosis(t)

⇒ hasPrognosis(r,t)

JCOHighGradeCaTam: Women(p) ∧ hasDisease(p,q) ∧ BreastCancer(q) ∧ hasDisease(p,r)

∧ EndometrialCa(r)

∧ hasTreatment(p,s) ∧ Tamoxifen(s)

∧ ReducedEndoCaDFS(t) ∧ refersDisease(t,r) ∧ hasValue(t,0.68)

⇒ hasDeltaRisk(r,t)

NSBABP-B14Tam: Women(p) ∧ hasDisease(p,q) ∧ BreastCancer(q) ∧ LNNeg(q) ∧ ER-Pos(q)

∧ hasTreatment(p,s) ∧ Tamoxifen5Yr(s)

∧ IncreasedRiskEndoCa(r) ∧ refersDisease(r,t) ∧ EndometrialCa(t) ∧ hasValue(r,7.5)

⇒ hasDeltaRisk(p,r)

LancetTamEndoCa: Women(p) ∧ hasDisease(p,q) ∧ BreastCancer(q)

∧ hasTreatment(p,s) ∧ Tamoxifen(s)

∧ IncreasedRiskEndoCa(r) ∧ refersDisease(r,s) ∧ (EndometrialCancer(s)) ∧ hasValue(r,2.3)

⇒ hasDeltaRisk(p,r)

GynaeOnc: Women(p) ∧ hasDisease(p,q) ∧ BreastCancer(q) ∧ hasDisease(p,r) ∧ EndometrialCa(r)

∧ hasTreatment(p,s) ∧ Tamoxifen(s)

∧ PoorPrognosis(t)

⇒ ¬ hasPrognosis(r,t)

RoswellPark: Women(p) ∧ hasDisease(p,q) ∧ BreastCancer(q)

∧ hasTreatment(p,s) ∧ Tamoxifen(s)

∧ RiskEndoCa(r) ∧ refersDisease(r,t) ∧ EndometrialCa(t) ∧ hasValue(r,0.005)

⇒ hasRisk(p,r)

Lancet94: Women(p)

∧ hasTreatment(p,s) ∧ Tamoxifen(s)

∧ RiskEndoAbnormality(r) ∧ refersDisease(r,s) ∧ EndoAbnormality(s) ∧ hasValue(r,3.9)

⇒ hasDeltaRisk(p,r)

BCRT98: Women(p)

∧ hasTreatment(p,s) ∧ Tamoxifen(s)

∧ RiskEndoAbnormality(r) ∧ refersDisease(r,s) ∧ EndoAbnormality(s) ∧ hasValue(r,0)

⇒ hasDeltaRisk(p,r)

Stockholm95: Women(p)

∧ hasDisease(p,q) ∧ EarlyBreastCancer(q) ∧ hasTreatment(p,s) ∧ TamoxifenHighDose(s)

∧ ReducedRiskBreastCancer(r) ∧ refersDisease(r,s) ∧ BreastCancer(s) ∧ hasValue(r,0.6)

⇒ hasDeltaRisk(p,r)

Stockholm95: Women(p)

∧ hasDisease(p,q) ∧ EarlyBreastCancer(q) ∧ hasTreatment(p,s) ∧ TamoxifenHighDose(s)

∧ RiskEndoCa(r) ∧ refersDisease(r,s) ∧ EndometrialCa(s) ∧ hasValue(r,6)

⇒ hasDeltaRisk(p,r)

Stockholm95: Women(p)

∧ hasDisease(p,q) ∧ EarlyBreastCancer(q) ∧ hasTreatment(p,s) ∧ TamoxifenHighDose(s)

∧ RiskGastricCa(r) ∧ refersDisease(r,s) ∧ GastricCancer(s) ∧ hasValue(r,3)

⇒ hasDeltaRisk(p,r)

StockholmJoint: Women(p)

∧ hasDisease(p,q) ∧ EarlyBreastCancer(q) ∧ hasTreatment(p,s) ∧ Tamoxifen(s)

∧ RiskEndoCa(r) ∧ refersDisease(r,t) ∧ EndometrialCa(t) ∧ hasValue(r,4.1)

⇒ hasDeltaRisk(p,r)

StockholmJoint: Women(p)

∧ hasDisease(p,q) ∧ EarlyBreastCancer(q) ∧ hasTreatment(p,s) ∧ Tamoxifen(s)

$\wedge$ RiskCRCa(r) $\wedge$ refersDisease(r,t) $\wedge$ ColorectalCa(t) $\wedge$ hasValue(r,1.9)

$\Rightarrow$ hasDeltaRisk(p,r)

StockholmJoint: Women(p)

$\wedge$ hasDisease(p,q) $\wedge$ EarlyBreastCancer(q) $\wedge$ hasTreatment(p,s) $\wedge$ Tamoxifen(s)

$\wedge$ RiskGastricCa(r) $\wedge$ refersDisease(r,s) $\wedge$ GastricCancer(s) $\wedge$ hasValue(r,3.2)

$\Rightarrow$ hasDeltaRisk(p,r)

ECOGVascular: Women(p)

$\wedge$ hasDisease(p,q) $\wedge$ BreastCancer(q) $\wedge$ hasTreatment(p,s) $\wedge$ AdjuvantTreatment(s)

$\wedge$ RiskVenousThrombus(r) $\wedge$ refersDisease(r,s) $\wedge$ VenousThrombus(s) $\wedge$ hasValue(r,3.4)

$\Rightarrow$ hasDeltaRisk(p,r)

ECOGVascular: Women(p)

$\wedge$ hasDisease(p,q) $\wedge$ BreastCancer(q) $\wedge$ hasTreatment(p,s) $\wedge$ Tamoxifen(s)

$\wedge$ RiskVenousThrombosis(r) $\wedge$ refersDisease(r,s) $\wedge$ VenousThrombus(s) $\wedge$ hasValue(r,3.4)

$\Rightarrow$ hasDeltaRisk(p,r)

ArchIntMedAntiThrombin: Women(p)

$\wedge$ hasDisease(p,q) $\wedge$ BreastCancer(q) $\wedge$ hasTreatment(p,s) $\wedge$ Tamoxifen(s) $\wedge$ PostMenopausal(p)

$\wedge$ ReducedPlasmaFibrinogenLevel(r) $\wedge$ refersSubstance(r,PlasmaFibrinogen) $\wedge$ hasValue(r,0.85)

$\Rightarrow$ hasDeltaLevel(p,r)

ArchIntMedAntiThrombin: Women(p)

$\wedge$ hasDisease(p,q) $\wedge$ BreastCancer(q) $\wedge$ hasTreatment(p,s) $\wedge$ Tamoxifen(s) $\wedge$ PostMenopausal(p)

$\wedge$ DecreasedPlasmaPlateletCount(r) $\wedge$ refersSubstance(r,PlateletLevel) $\wedge$ hasValue(r,0.92)

$\Rightarrow$ hasDeltaLevel(p,r)

JNCITamPrev: Women(p)

$\wedge$ hasTreatment(p,s) $\wedge$ Tamoxifen(s)

$\wedge$ RiskCataracts(r) $\wedge$ refersDisease(r,ProtoCataract) $\wedge$ hasValue(r,0.14)

$\Rightarrow$ hasDeltaRisk(p,r)

JNCITamPrev: Women(p)

$\wedge$ hasTreatment(p,s) $\wedge$ Tamoxifen(s)

$\wedge$ IncreasedRiskEndoCa(r) $\wedge$ refersDisease(r,t) $\wedge$ EndometrialCa(t) $\wedge$ hasValue(r,2.53)

$\Rightarrow$ hasDeltaRisk(p,r)

AmJObsGynae1996: Women(p)

$\wedge$ hasTreatment(p,s) $\wedge$ Tamoxifen20mg(s)

$\wedge$ RiskOvarianCyst(r) $\wedge$ refersDisease(r, ProtoOvarianCysts) $\wedge$ hasValue(r,0.11)

$\Rightarrow$ hasRisk(p,r)

GynaeOncol1996: Women(p)

$\wedge$ hasTreatment(p,s) $\wedge$ Tamoxifen20mg(s)

$\wedge$ RiskOvarianCyst(r) $\wedge$ refersDisease(r, ProtoOvarianCysts) $\wedge$ hasValue(r,0.06)

$\Rightarrow$ hasRisk(p,r)

ArchIntMed1991: Women(p) $\wedge$ Postmenopausal(p)

$\wedge$ hasTreatment(p,s) $\wedge$ Tamoxifen(s)

$\wedge$ IncreasedRiskPersistentSideEffects(r) $\wedge$ refersDisease(r,t) $\wedge$ PersistentSideEffects(t) $\wedge$ hasValue(r,2.3)

$\Rightarrow$ hasDeltaRisk(p,r)

JCO1994Clonidine: Women(p) $\wedge$ hasSymptoms(p,q) $\wedge$ HotFlushes(q)

$\wedge$ hasDisease(p,u) $\wedge$ BreastCancer(u) $\wedge$ hasTreatment(p,s) $\wedge$ Clonidine(s)

$\wedge$ ReducedRiskHotFlushes(r) $\wedge$ refersDisease(r,q) $\wedge$ hasValue(r,0.8)

$\Rightarrow$ hasDeltaRisk(p,r)

BMJ1992: Women(p)

$\wedge$ hasDisease(p,u) $\wedge$ BreastCancer(u) $\wedge$ hasTreatment(p,s) $\wedge$ Tamoxifen(s)

$\wedge$ IncreasedRiskRetinopathy(r) $\wedge$ refersDisease(r,q) $\wedge$ Retinopathy(q)

$\Rightarrow$ hasDeltaRisk(p,r)

Eye1999: Women(p)

$\wedge$ hasDisease(p,u) $\wedge$ BreastCancer(u) $\wedge$ hasTreatment(p,s) $\wedge$ Tamoxifen(s)

$\wedge$ RiskRetinopathy(r) $\wedge$ hasValue(r,0.12) $\wedge$ refersDisease(r,q) $\wedge$ Retinopathy(q)

$\Rightarrow$ hasRisk(p,r)

ArchIntMed1991: Women(p)

∧ hasDisease(p,u) ∧ NodeNegativeBreastCancer(u) ∧ hasTreatment(p,s) ∧ Tamoxifen(s)

∧ ReducedPlasmaCholesterol(r) ∧ hasValue(r,0.88) ∧ refersSubstance(r,q) ∧ PlasmaCholesterol(q)

⇒ hasDeltaLevel(p,r)

ArchIntMed1991: Women(p)

∧ hasDisease(p,u) ∧ NodeNegativeBreastCancer(u) ∧ hasTreatment(p,s) ∧ Tamoxifen(s)

∧ ReducedLDL(r) ∧ hasValue(r,0.8) ∧ refersSubstance(r,q) ∧ PlasmaLDL(q)

⇒ hasDeltaLevel(p,r)

BMJ1991: Women(p)

∧ hasDisease(p,u) ∧ BreastCancer(u) ∧ hasTreatment(p,s) ∧ Tamoxifen(s)

∧ ReducedRiskFatalMI(r) ∧ hasValue(r,0.37) ∧ refersDisease(r,q) ∧ FatalMI(q)

⇒ hasDeltaRisk(p,r)

JNCI1993: Women(p)

∧ hasDisease(p,u) ∧ BreastCancer(u) ∧ hasTreatment(p,s) ∧ Tamoxifen(s)

∧ ReducedRiskCardiacDisease(r) ∧ hasValue(r,0.68) ∧ refersDisease(r,q) ∧ CardiacDisease(q)

⇒ hasDeltaRisk(p,r)

JNCI1997: Women(p)

∧ hasDisease(p,u) ∧ BreastCancer(u) ∧ hasTreatment(p,s) ∧ Tamoxifen(s)

∧ IncreasedDiseaseSpecificSurvival(r) ∧ hasValue(r,1.52) ∧ refersDisease(r,q) ∧ CardiacDisease(q)

⇒ hasDeltaRisk(p,r)

NEJM1992: Women(p) ∧ Postmenopausal(p)

∧ hasDisease(p,u) ∧ NodeNegativeBreastCancer(u) ∧ hasTreatment(p,s) ∧ Tamoxifen(s)

∧ IncreasedLumbarBMD(r) ∧ hasValue(r,1.006) ∧ refersSubstance(r,q) ∧ LumbarBMD(q)

⇒ hasDeltaLevel(p,r)

JCOTamBone1994: Women(p) ∧ Postmenopausal(p)

∧ hasDisease(p,u) ∧ NodeNegativeBreastCancer(u) ∧ hasTreatment(p,s) ∧ Tamoxifen30mg(s)

$\wedge$ IncreasedLumbarBMD(r) $\wedge$ refersSubstance(r,q) $\wedge$ LumbarBMD(q)

$\Rightarrow$ hasDeltaLevel(p,r)

JCOTamBone1994: Women(p) $\wedge$ Postmenopausal(p)

$\wedge$ hasDisease(p,u) $\wedge$ NodeNegativeBreastCancer(u)

$\wedge$ DecreasedRadialBMD(r) $\wedge$ refersSubstance(r,q) $\wedge$ RadialBMD(q)

$\Rightarrow$ hasDeltaLevel(p,r)

JCOTamBone1994: Women(p) $\wedge$ Postmenopausal(p)

$\wedge$ hasDisease(p,u) $\wedge$ NodeNegativeBreastCancer(u) $\wedge$ hasTreatment(p,s) $\wedge$ Tamoxifen30mg(s)

$\wedge$ UnchangedRadialBMD(r) $\wedge$ refersSubstance(r,q) $\wedge$ RadialBMD(q)

$\Rightarrow$ hasDeltaLevel(p,r)

ArchIntMed1994: Women(p) $\wedge$ Postmenopausal(p)

$\wedge$ hasDisease(p,u) $\wedge$ NodeNegativeBreastCancer(u) $\wedge$ hasTreatment(p,s) $\wedge$ Tamoxifen5Yr(s)

$\wedge$ IncreasedLumbarBMD(r) $\wedge$ hasValue(r,1.008) $\wedge$ refersSubstance(r,q) $\wedge$ LumbarBMD(q)

$\Rightarrow$ hasDeltaLevel(p,r)

ArchIntMed1994: Women(p) $\wedge$ Postmenopausal(p)

$\wedge$ hasDisease(p,u) $\wedge$ NodeNegativeBreastCancer(u)

$\wedge$ DecreasedLumbarBMD(r) $\wedge$ hasValue(r,0.993) $\wedge$ refersSubstance(r,q) $\wedge$ LumbarBMD(q)

$\Rightarrow$ hasDeltaLevel(p,r)

JCO1996: Women(p) $\wedge$ Postmenopausal(p)

$\wedge$ hasTreatment(p,q) $\wedge$ Tamoxifen20mg(q)

$\wedge$ IncreasedLumbarBMD(r) $\wedge$ hasValue(r,1.012) $\wedge$ refersSubstance(r,q) $\wedge$ LumbarBMD(q)

$\Rightarrow$ hasDeltaLevel(p,r)

JC01996: Women(p) $\wedge$ Premenopausal(p)

$\wedge$ hasTreatment(p,q) $\wedge$ Tamoxifen20mg(q)

$\wedge$ DecreasedLumbarBMD(r) $\wedge$ hasValue(r,0.986) $\wedge$ refersSubstance(r,q) $\wedge$ LumbarBMD(q)

$\Rightarrow$ hasDeltaLevel(p,r)

The application of $Rewrite()$ to these rules, in conjunction with the values given in Chapter 7 produced the following rules:

IBCSGT13-93a: Women(p) $\wedge$ hasDisease(p,q) $\wedge$ BreastCancer(q) $\wedge$ LNPos(q) $\wedge$ ERPos(q)

$\wedge$ Premenopausal(p)

$\wedge$ hasPosValue(p,t) $\wedge$ Tamoxifen5YrCourse(s)

$\wedge$ IncreasedBrCaDFS(t)

$\Rightarrow$ hasPosIntent(p,s)

IBCSGT13-93a: Women(p) $\wedge$ hasDisease(p,q) $\wedge$ BreastCancer(q) $\wedge$ LNPos(q) $\wedge$ ERNeg(q)

$\wedge$ Premenopausal(p)

$\wedge$ hasPosValue(p,t) $\wedge$ Tamoxifen5YrCourse(s)

$\wedge$ DecreasedBrCaDFS(t)

$\Rightarrow$ hasPosIntent(p,s)

IBCSGT13-93a: Women(p) $\wedge$ hasDisease(p,q) $\wedge$ BreastCancer(q) $\wedge$ LNPos(q) $\wedge$ ERPos(q)

$\wedge$ Premenopausal(p) $\wedge$ Amenorrhic(p)

$\wedge$ hasPosValue(p,t) $\wedge$ Tamoxifen5YrCourse(s)

$\wedge$ IncreasedBrCaDFS(t)

$\Rightarrow$ hasPosIntent(p,s)

NSABP14a: Women(p) $\wedge$ hasDisease(p,q) $\wedge$ BreastCancer(q) $\wedge$ LNNeg(q)

$\wedge$ hasNegValue(p,t) $\wedge$ TamoxifenMoreThan5YrCourse(s)

$\wedge$ DecreasedBrCaDFS(t)

$\Rightarrow$ hasNegIntent(p,s)

NASBP-B16: Women(p) $\wedge$ hasDisease(p,q) $\wedge$ BreastCancer(q) $\wedge$ Age50Plus(p)

$\wedge$ hasPosValue(p,t) $\wedge$ ACTChemoTam(s)

$\wedge$ IncreasedBrCaDFS(t)

$\Rightarrow$ hasPosIntent(p,s)

NASBP-B16: Women(p) $\wedge$ hasDisease(p,q) $\wedge$ BreastCancer(q) $\wedge$ Age50Plus(p)

$\wedge$ hasPosValue(p,t) $\wedge$ ACTChemoTam(s)

$\wedge$ IncreasedOS(t)

$\Rightarrow$ hasPosIntent(p,s)

NASBP-B16: Women(p) $\wedge$ hasDisease(p,q) $\wedge$ BreastCancer(q) $\wedge$ Age50Plus(p)

$\wedge$ hasTreatment(p,t) $\wedge$ PAFTChemoTam(s)

$\wedge$ IncreasedBrCaDFS(t)

$\Rightarrow$ hasPosIntent(p,s)

NASBP-B16: Women(p) $\wedge$ hasDisease(p,q) $\wedge$ BreastCancer(q) $\wedge$ Age50Plus(p)

$\wedge$ hasPosValue(p,s) $\wedge$ PFTChemoTam(s)

$\wedge$ IncreasedBrCaDFS(t)

$\Rightarrow$ hasPosIntent(p,s)

NSABP20a: Women(p) $\wedge$ hasDisease(p,q) $\wedge$ BreastCancer(q) $\wedge$ ERPos(q) $\wedge$ LNNeg(q)

$\wedge$ hasPosValue(p,t) $\wedge$ MFTChemoTam(s)

$\wedge$ IncreasedBrCaDFS(t)

$\Rightarrow$ hasPosIntent(p,s)

NSABP20a: Women(p) $\wedge$ hasDisease(p,q) $\wedge$ BreastCancer(q) $\wedge$ ERPos(q) $\wedge$ LNNeg(q)

$\wedge$ hasPosValue(p,s) $\wedge$ CMFTChemoTam(s)

$\wedge$ IncreasedBrCaDFS(t)

$\Rightarrow$ hasPosIntent(p,s)

NSABP20a: Women(p) $\wedge$ hasDisease(p,q) $\wedge$ BreastCancer(q) $\wedge$ ERPos(q) $\wedge$ LNNeg(q)

$\wedge$ hasPosValue(p,t) $\wedge$ MFTChemoTam(s)

$\wedge$ IncreasedOS(t)

$\Rightarrow$ hasPosIntent(p,s)

NSABP20a: Women(p) $\wedge$ hasDisease(p,q) $\wedge$ BreastCancer(q) $\wedge$ ERPos(q) $\wedge$ LNNeg(q)

$\wedge$ hasPosValue(p,t) $\wedge$ CMFTChemoTam(s)

$\wedge$ IncreasedOS(t)

$\Rightarrow$ hasPosIntent(p,s)

ICBSG97a: Women(p) $\wedge$ hasDisease(p,q) $\wedge$ BreastCancer(q) $\wedge$ Postmenopausal(p) $\wedge$ LNPos(q)

$\wedge$ hasPosValue(p,t) $\wedge$ EarlyCMFTChemoTam(s)

$\wedge$ IncreasedBrCaDFS(t)

$\Rightarrow$ hasPosIntent(p,s)

ICBSG97a: Women(p) $\wedge$ hasDisease(p,q) $\wedge$ BreastCancer(q) $\wedge$ Postmenopausal(p) $\wedge$ LNPos(q)

$\wedge$ hasPosValue(p,t) $\wedge$ LateCMFTChemoTam(s)

∧ NoChangeBrCaDFS(t)

⇒ hasPosIntent(p,s)

ICBSG97a: Women(p) ∧ hasDisease(p,q) ∧ BreastCancer(q) ∧ Postmenopausal(p) ∧ LNPos(q) ∧ ERPos(q)

∧ hasPosValue(p,t) ∧ CMFTChemoTam(s)

∧ IncreasedBrCaDFS(t)

⇒ hasPosIntent(p,s)

NCICanada: Women(p) ∧ hasDisease(p,q) ∧ BreastCancer(q) ∧ Postmenopausal(p) ∧ LNPos(q) ∧ HRPos(q)

∧ hasPosValue(p,t) ∧ CMFTChemoTam(s)

∧ IncreasedSideEffects(t)

⇒ hasPosIntent(p,s)

Lancet89: Women(p) ∧ hasDisease(p,q) ∧ BreastCancer(q) ∧ Postmenopausal(p)

∧ hasPosValue(p,t) ∧ Tamoxifen(s)

∧ ReducedRiskBrCa(t) ∧ hasValue(t, 0.55)

⇒ hasPosIntent(p,s)

Lancet89: Women(p) ∧ hasDisease(p,q) ∧ BreastCancer(q) ∧ Postmenopausal(p)

∧ hasNegValue(p,t) ∧ Tamoxifen(s)

∧ IncreasedRiskEndoCa(t) ∧ hasValue(t, 6.4)

⇒ hasNegIntent(p,s)

JCOHighGradeCaTam: Women(p) ∧ hasDisease(p,q) ∧ BreastCancer(q) ∧ hasDisease(p,r)

∧ EndometrialCa(r) hasPrognosis(r,t)

∧ hasNegValue(p,r) ∧ Tamoxifen(s)

∧

⇒ hasNegIntent(p,s)

JCOHighGradeCaTam: Women(p) ∧ hasDisease(p,q) ∧ BreastCancer(q) ∧ hasDisease(p,r)

∧ EndometrialCa(r)

∧ hasNegValue(p,t) ∧ Tamoxifen(s)

∧ ReducedEndoCaDSS(t)

⇒ hasNegIntent(r,s)

NSBABP-B14Tam: Women(p) ∧ hasDisease(p,q) ∧ BreastCancer(q) ∧ LNNeg(q) ∧ ER-Pos(q)

∧ hasNegValue(p,r) ∧ Tamoxifen5Yr(s)

∧ IncreasedRiskEndoCa(r)

⇒ hasNegIntent(p,s)

LancetTamEndoCa: Women(p) ∧ hasDisease(p,q) ∧ BreastCancer(q)

∧ hasNegValue(p,r) ∧ Tamoxifen(s)

∧ IncreasedRiskEndoCa(r)

⇒ hasNegIntent(p,s)

GynaeOnc: Women(p) ∧ hasDisease(p,q) ∧ BreastCancer(q) ∧ hasDisease(p,r) ∧ EndometrialCa(r)

∧ hasTreatment(p,s) ∧ Tamoxifen(s)

∧ PoorPrognosis(t)

⇒ ¬ hasPrognosis(r,t)

Lancet94: Women(p)

∧ hasNegValue(p,r) ∧ Tamoxifen(s)

∧ RiskEndoAbnormality(r)

⇒ hasNegIntent(p,s)

Stockholm95: Women(p)

∧ hasDisease(p,q) ∧ EarlyBreastCancer(q) ∧ hasTreatment(p,s) ∧ TamoxifenHighDose(s)

∧ RiskBreastCancer(r) ∧ hasValue(r,0.6)

⇒ hasDeltaRisk(p,r)

Stockholm95: Women(p)

∧ hasDisease(p,q) ∧ EarlyBreastCancer(q) ∧ hasNegValue(p,r) ∧ TamoxifenHighDose(s)

∧ IncreasedRiskEndoCa(r) ∧ hasValue(r,6)

⇒ hasNegIntent(p,s)

Stockholm95: Women(p)

∧ hasDisease(p,q) ∧ EarlyBreastCancer(q) ∧ hasNegValue(p,r) ∧ TamoxifenHighDose(s)

∧ IncreasedRiskGastricCa(r) ∧ hasValue(r,3)

⇒ hasNegIntent(p,s)

StockholmJoint: Women(p)

$\wedge$ hasDisease(p,q) $\wedge$ EarlyBreastCancer(q) $\wedge$ hasNegValue(p,r) $\wedge$ Tamoxifen(s)

$\wedge$ IncreasedRiskEndoCa(r) $\wedge$ hasValue(r,4.1)

$\Rightarrow$ hasNegIntent(p,s)

StockholmJoint: Women(p)

$\wedge$ hasDisease(p,q) $\wedge$ EarlyBreastCancer(q) $\wedge$ hasNegValue(p,r) $\wedge$ Tamoxifen(s)

$\wedge$ IncreasedRiskCRCa(r) $\wedge$ hasValue(r,1.9)

$\Rightarrow$ hasNegIntent(p,s)

StockholmJoint: Women(p)

$\wedge$ hasDisease(p,q) $\wedge$ EarlyBreastCancer(q) $\wedge$ hasNegValue(p,r) $\wedge$ Tamoxifen(s)

$\wedge$ IncreasedRiskGastricCa(r) $\wedge$ hasValue(r,3.2)

$\Rightarrow$ hasNegIntent(p,s)

ECOGVascular: Women(p)

$\wedge$ hasDisease(p,q) $\wedge$ BreastCancer(q) $\wedge$ hasNegValue(p,r) $\wedge$ AdjuvantTreatment(s)

$\wedge$ IncreasedRiskVenousThrombus(r)

$\Rightarrow$ hasNegIntent(p,s)

ECOGVascular: Women(p)

$\wedge$ hasDisease(p,q) $\wedge$ BreastCancer(q) $\wedge$ hasNegValue(p,r) $\wedge$ Tamoxifen(s)

$\wedge$ IncreasedRiskVenousThrombosis(r)

$\Rightarrow$ hasNegIntent(p,s)

JNCITamPrev: Women(p)

$\wedge$ hasNegValue(p,s) $\wedge$ Tamoxifen(s)

$\wedge$ RiskCataracts(r)

$\Rightarrow$ hasNegIntent(p,s)

JNCITamPrev: Women(p)

$\wedge$ hasNegValue(p,r) $\wedge$ Tamoxifen(s)

$\wedge$ IncreasedRiskEndoCa(r)

$\Rightarrow$ hasNegIntent(p,s)

AmJObsGynae1996: Women(p)

$\wedge$ hasNegValue(p,r) $\wedge$ Tamoxifen20mg(s)

∧ RiskOvarianCyst(r)

⇒ hasNegIntent(p,s)

ArchIntMed1991: Women(p) ∧ Postmenopausal(p)

∧ hasNegValue(p,r) ∧ Tamoxifen(s)

∧ IncreasedRiskPersistentSideEffects(r)

⇒ hasNegIntent(p,s)

JCO1994: Women(p) ∧ hasSymptoms(p,q) ∧ HotFlushes(q)

∧ hasDisease(p,u) ∧ BreastCancer(u) ∧ hasPosValue(p,r) ∧ Clonidine(s)

∧ ReducedRiskHotFlushes(r)

⇒ hasPosIntent(p,s)

BMJ1992: Women(p)

∧ hasDisease(p,u) ∧ BreastCancer(u) ∧ hasNegValue(p,r) ∧ Tamoxifen(s)

∧ IncreasedRiskRetinopathy(r)

⇒ hasNegIntent(p,s)

ArchIntMed1991: Women(p)

∧ hasDisease(p,u) ∧ BreastCancer(u) ∧ LNNeg(U) ∧ hasTreatment(p,s) ∧ Tamoxifen(s)

∧ hasPosValue(p,r) ∧ ReducedPlasmaCholesterol(r)

⇒ hasPosIntent(p,s)

ArchIntMed1991: Women(p)

∧ hasDisease(p,u) ∧ BreastCancer(u) ∧ LNNeg(U) ∧ hasTreatment(p,s) ∧ Tamoxifen(s)

∧ hasPosIntent(p,r) ∧ ReducedLDL(r)

⇒ hasPosIntent(p,s)

BMJ1991: Women(p)

∧ hasDisease(p,u) ∧ BreastCancer(u) ∧ hasPosValue(p,r) ∧ Tamoxifen(s)

∧ RecucedRiskFatalMI(r)

⇒ hasPosIntent(p,s)

JNCI1993: Women(p)

∧ hasDisease(p,u) ∧ BreastCancer(u) ∧ hasPosValue(p,r) ∧ Tamoxifen(s)

∧ ReducedRiskCardiacDisease(r)

$\Rightarrow$ hasPosIntent(p,s)

JNCI1997: Women(p)

$\wedge$ hasDisease(p,u) $\wedge$ BreastCancer(u) $\wedge$ hasPosValue(p,r) $\wedge$ Tamoxifen(s)

$\wedge$ IncreasedDiseaseSpecificSurvival(r)

$\Rightarrow$ hasPosIntent(p,s)

NEJM1992: Women(p) $\wedge$ Postmenopausal(p)

$\wedge$ hasDisease(p,u) $\wedge$ BreastCancer(u) $\wedge$ LNNeg(U) $\wedge$ hasTreatment(p,s) $\wedge$ Tamoxifen(s)

$\wedge$ hasPosValue(p,r) $\wedge$ IncreasedLumbarBMD(r)

$\Rightarrow$ hasPosIntent(p,s)

JCOTamBone1994: Women(p) $\wedge$ Postmenopausal(p)

$\wedge$ hasDisease(p,u) $\wedge$ BreastCancer(u) $\wedge$ LNNeg(U) $\wedge$ hasTreatment(p,s) $\wedge$ Tamoxifen30mg(s)

$\wedge$ hasPosValue(p,r) $\wedge$ IncreasedLumbarBMD(r)

$\Rightarrow$ hasPosIntent(p,s)

JCOTamBone1994: Women(p) $\wedge$ Postmenopausal(p)

$\wedge$ hasDisease(p,u) $\wedge$ BreastCancer(u) $\wedge$ LNNeg(U)

$\wedge$ hasNegValue(p,r) $\wedge$ DecreasedRadialBMD(r)

$\Rightarrow$ hasNegIntent(p,s)

ArchIntMed1994: Women(p) $\wedge$ Postmenopausal(p)

$\wedge$ hasDisease(p,u) $\wedge$ NodeNegativeBreastCancer(u) $\wedge$ Tamoxifen5Yr(s)

$\wedge$ hasPosValue(p,r) $\wedge$ IncreasedLumbarBMD(r)

$\Rightarrow$ hasPosIntent(p,s)

ArchIntMed1994: Women(p) $\wedge$ Postmenopausal(p)

$\wedge$ hasDisease(p,u) $\wedge$ NodeNegativeBreastCancer(u) $\wedge$ Tamoxifen5Yr(s)

$\wedge$ hasNegValue(p,r) $\wedge$ DecreasedLumbarBMD(r)

$\Rightarrow$ hasNegIntent(p,s)

JCO1996: Women(p) $\wedge$ Postmenopausal(p)

$\wedge$ Tamoxifen(s)

$\wedge$ hasPosValue(p,r) $\wedge$ IncreasedLumbarBMD(r)

$\Rightarrow$ hasPosIntent(p,s)

JC01996: Women(p) $\wedge$ Premenopausal(p)

$\wedge$ Tamoxifen(s)

$\wedge$ hasNegValue(p,r) $\wedge$ IncreasedLumbarBMD(r)

$\Rightarrow$ hasNegIntent(p,s)

Appendix C: Sources

The information on the source of each rules is as follows

Source: **EBCTCG2005a**

$\wedge$ hasDate(EBCTCG2005a,2005)

$\wedge$ hasEvidenceLevel(EBCTCG2005a,1a)

$\wedge$ hasNumber(EBCTCG2005a,8311)

$\wedge$ hasArea(EBCTCG2005a,World)

Source: **EBCTCG2005b**

$\wedge$ hasDate(EBCTCG2005b,2005)

$\wedge$ hasEvidenceLevel(EBCTCG2005b,1a)

$\wedge$ hasNumber(EBCTCG2005b,14336)

$\wedge$ hasArea(EBCTCG2005b,World)

Source: **IBCSGT13-93a**

$\wedge$ hasDate(IBCSGT13-93a,1993)

$\wedge$ hasEvidenceLevel(IBCSGT13-93a,1b)

$\wedge$ hasNumber(IBCSGT13-93a,1246)

$\wedge$ hasArea(IBCSGT13-93a,NonUSWorld)

Source: **NSABP14a**

$\wedge$ hasDate(NSABP14a,2001)

$\wedge$ hasEvidenceLevel(NSABP14a,1b)

$\wedge$ hasNumber(NSABP14a,1172)

$\wedge$ hasArea(NSABP14a,USA)

Source: **ScotTamTrial**

$\wedge$ hasDate(ScotTamTrial,2001)

$\wedge$ hasEvidenceLevel(ScotTamTrial,1b)

$\wedge$ hasNumber(ScotTamTrial,1323)

$\wedge$ hasArea(ScotTamTrial,Scotland)

Source: **ECOGa**

$\wedge$ hasDate(ECOGa,1996)

∧ hasEvidenceLevel(ECOGa,1b)

∧ hasNumber(ECOGa,193)

∧ hasArea(ECOGa,USA)

Source: **JNCIEd**

∧ hasDate(JNCIEd,1996)

∧ hasEvidenceLevel(JNCIEd,5)

∧ hasNumber(JNCIEd,0)

∧ hasArea(JNCIEd,World)

Source: **NASBP-B16**

∧ hasDate(NASBP-B16,1990)

∧ hasEvidenceLevel(NASBP-B16,1b)

∧ hasNumber(NASBP-B16,1124)

∧ hasArea(NASBP-B16,USA)

Source: **NSABP20a**

∧ hasDate(NSABP20a,1997)

∧ hasEvidenceLevel(NSABP20a,1b)

∧ hasNumber(NSABP20a,2306)

∧ hasArea(NSABP20a,USA)

Source: **IBCSG97a**

∧ hasDate(IBCSG97a,1990)

∧ hasEvidenceLevel(IBCSG97a,1b)

∧ hasNumber(IBCSG97a,1266)

∧ hasArea(IBCSG97a,World)

Source: **NCICanada**

∧ hasDate(NCICanada,1997)

∧ hasEvidenceLevel(NCICanada,1b)

∧ hasNumber(NCICanada,705)

∧ hasArea(NCICanada,Canada)

Source: **Lancet89**

∧ hasDate(Lancet89,1989)

∧ hasEvidenceLevel(Lancet89,2a)

∧ hasNumber(Lancet89,1846)

∧ hasArea(Lancet89,Sweden)

Source: **NSBABP-B14Tam**

∧ hasDate(NSBABP-B14Tam,1994)

∧ hasEvidenceLevel(NSBABP-B14Tam,1b)

$\wedge$ hasNumber(NSBABP-B14Tam,4063)

$\wedge$ hasArea(NSBABP-B14Tam,USA)

Source: **LancetTamEndoCa**

$\wedge$ hasDate(LancetTamEndoCa,1994)

$\wedge$ hasEvidenceLevel(LancetTamEndoCa,4)

$\wedge$ hasNumber(LancetTamEndoCa,383)

$\wedge$ hasArea(LancetTamEndoCa,Netherlands)

Source: **GynaeOnc**

$\wedge$ hasDate(GynaeOnc,1994)

$\wedge$ hasEvidenceLevel(GynaeOnc,4)

$\wedge$ hasNumber(GynaeOnc,73)

$\wedge$ hasArea(GynaeOnc,USA)

Source: **RoswellPark**

$\wedge$ hasDate(RoswellPark,1996)

$\wedge$ hasEvidenceLevel(RoswellPark,4)

$\wedge$ hasNumber(RoswellPark,4597)

$\wedge$ hasArea(RoswellPark,USA)

Source: **Lancet94**

$\wedge$ hasDate(Lancet94,1994)

$\wedge$ hasEvidenceLevel(Lancet94,1b)

$\wedge$ hasNumber(Lancet94,111)

$\wedge$ hasArea(Lancet94,UK)

Source: **BCRT98**

$\wedge$ hasDate(BCRT98,1998)

$\wedge$ hasEvidenceLevel(BCRT98,4)

$\wedge$ hasNumber(BCRT98,164)

$\wedge$ hasArea(BCRT98,Italy)

Source: **StockholmJoint**

$\wedge$ hasDate(StockholmJoint,1995)

$\wedge$ hasEvidenceLevel(StockholmJoint,4)

$\wedge$ hasNumber(StockholmJoint,4914)

$\wedge$ hasArea(StockholmJoint,Scandinavia)

Source: **Stockholm95**

$\wedge$ hasDate(StockholmJoint,1995)

$\wedge$ hasEvidenceLevel(StockholmJoint,4)

$\wedge$ hasNumber(StockholmJoint,2729)

$\wedge$ hasArea(StockholmJoint,Sweden)

Source: **ECOGVascular**

∧ hasDate(ECOGVascular,1991)

∧ hasEvidenceLevel(ECOGVascular,2b)

∧ hasNumber(ECOGVascular,2673)

∧ hasArea(ECOGVascular,USA)

Source: **ArchIntMedAntiThrombin**

∧ hasDate(ArchIntMedAntiThrombin,1992)

∧ hasEvidenceLevel(ArchIntMedAntiThrombin,2a)

∧ hasNumber(ArchIntMedAntiThrombin,140)

∧ hasArea(ArchIntMedAntiThrombin,USA)

Source: **JNCITamPrev**

∧ hasDate(JNCITamPrev,1998)

∧ hasEvidenceLevel(JNCITamPrev,1b)

∧ hasNumber(JNCITamPrev,13388)

∧ hasArea(JNCITamPrev,USA)

Source: **JCOHighGradeCaTam**

∧ hasDate(JCOHighGradeCaTam,1993)

∧ hasEvidenceLevel(JCOHighGradeCaTam,4)

∧ hasNumber(JCOHighGradeCaTam,53)

∧ hasArea(JCOHighGradeCaTam,USA)

Source: **AmJObsGynae1996**

∧ hasDate(AmJObsGynae1996,1996)

∧ hasEvidenceLevel(AmJObsGynae1996,4)

∧ hasNumber(AmJObsGynae1996,95)

∧ hasArea(AmJObsGynae1996,Israel)

Source: **GynaeOncol1996**

∧ hasDate(GynaeOncol1996,1996)

∧ hasEvidenceLevel(GynaeOncol1996,4)

∧ hasNumber(GynaeOncol1996,175)

∧ hasArea(GynaeOncol1996,Israel)

Source: **ArchIntMed1991**

∧ hasDate(ArchIntMed1991,1991)

∧ hasEvidenceLevel(ArchIntMed1991,1c)

∧ hasNumber(ArchIntMed1991,140)

∧ hasArea(ArchIntMed1991,USA)

Source: **JCO1994Clondine**

∧ hasDate(JCO1994,1994)

∧ hasEvidenceLevel(JCO1994,1b)

∧ hasNumber(JCO1994,-1)

∧ hasArea(JCO1994,USA)

Source: **BMJ1992**

∧ hasDate(BMJ1992,1992)

∧ hasEvidenceLevel(BMJ1992,5)

∧ hasNumber(BMJ1992,1)

∧ hasArea(BMJ1992,UK)

Source: **Eye1999**

∧ hasDate(Eye1999,1999)

∧ hasEvidenceLevel(Eye1999,1c)

∧ hasNumber(Eye1999,65)

∧ hasArea(Eye1999,Lebanon)

Source: **ArchIntMed1991**

∧ hasDate(ArchIntMed1991,1991)

∧ hasEvidenceLevel(ArchIntMed1991,1b)

∧ hasNumber(ArchIntMed1991,140)

∧ hasArea(ArchIntMed1991,USA)

Source: **BMJ1991**

∧ hasDate(BMJ1991,1991)

∧ hasEvidenceLevel(,2b)

∧ hasNumber(BMJ1991,1070)

∧ hasArea(BMJ1991,Scotland)

Source: **JNCI1993**

∧ hasDate(JNCI1993,1993)

∧ hasEvidenceLevel(JNCI1993,2b)

∧ hasNumber(JNCI1993,2363)

∧ hasArea(JNCI1993,Sweden)

Source: **JNCI1997**

∧ hasDate(JNCI1997,1997)

∧ hasEvidenceLevel(JNCI1997,2b)

∧ hasNumber(JNCI1997,2885)

∧ hasArea(JNCI1997,USA)

Source: **NEJM1992**

∧ hasDate(NEJM1992,1992)

∧ hasEvidenceLevel(NEJM1992,1b)

$\wedge$ hasNumber(NEJM1992,140)

$\wedge$ hasArea(NEJM1992,USA)

Source: **JCOTamBone1994**

$\wedge$ hasDate(JCOTamBone1994,1994)

$\wedge$ hasEvidenceLevel(JCOTamBone1994,1b)

$\wedge$ hasNumber(JCOTamBone1994,43)

$\wedge$ hasArea(JCOTamBone1994,Denmark)

Source: **ArchIntMed1994**

$\wedge$ hasDate(ArchIntMed1994,1994)

$\wedge$ hasEvidenceLevel(ArchIntMed1994,1b)

$\wedge$ hasNumber(ArchIntMed1994,140)

$\wedge$ hasArea(ArchIntMed1994,USA)

Source: **JCO1996**

$\wedge$ hasDate(JCO1996,1996)

$\wedge$ hasEvidenceLevel(JCO1996,1b)

$\wedge$ hasNumber(JCO1996,179)

$\wedge$ hasArea(JCO1996,UK)

Appendix D: Electronic Materials

Attached to the thesis is a CD. On this, I have provided the Breast Cancer ontology, Δ_{BC}, in an electronic format. It is provided as an OWL 2 file, and can be viewed either in a text editor or using a suitable ontology viewing and editing tool. Options include Protege (recommended) or Swoop, both of which require Java, or the OwlSight webservice which avoids downloads.

I also provide the prototype implementation, written in Java. It requires Java 1.5 or higher to run. A suitable Java VM is available free from Sun and other providers. For instructions on use, please see the README.txt file.

Should you have any problems with these files, please contact me directly. All files are available for download from http://www.acl.icnet.uk/~mw

Index

Bibliography

[1] R. L. Ackoff. From data to wisdom. *Journal of Applied Systems Analysis*, 16:3–9., 1989.

[2] AJCC. *AJCC Cancer Staging Manual*. Lippincott Raven, USA, 1997.

[3] L. Amgoud. A general argumentation framework for inference and decision making. In F. Bacchus and T. Jaakkola, editors, *In 21st Conference on Uncertainty in Artificial Intelligence*, pages 26–33. AUAI Press, July 2005.

[4] L. Amgoud. A unified setting for inference and decision: An argumentation-based approach. In C. Reed, editor, *5th Workshop on Computational Models of Natural Argument, CMNA 05.*, 2005.

[5] L. Amgoud, L. Bodenstaff, M. Caminada, P. McBurney, S. Parsons, H. Prakken, J. van Veenen, and G. Vreeswijk. Aspic - final review and report on formal argumentation system. Technical report, Universite Paul Sabatier (UPS), 2006.

[6] L. Amgoud, C. Cayrol, and M. Lagasquie-Schiex. On Bipolarity in Argumentation frameworks. In *Non-Monotonic Reasoning*, pages 1–9, 2004.

[7] G. Antoniou, A. Bikakis, and G. Wagner. A Defeasible Logic Programming System for the Web. In *Proceedings of 16th IEEE International Conference on Tools with Artificial Intelligence (ICTAI'04)*. IEEE, IEEE Press, 2004.

[8] K. Atkinson, T. Bench-Capon, and P. McBurney. Computational representation of practical argument. *Knowledge, Rationality and Action, a special section of Synthese*, 152(2):157–206, 2006.

[9] F. Baader, D. Calvanese, D. McGuinness, D. Nardi, and P. Patel-Schneider, editors. *The Description Logic Handbook: Theory, Implementation and Applications*. Cambridge University Press, January 2003.

[10] T. Bench-Capon and H. Prakken. Justifying actions by accruing arguments. In P. Dunne and T. Bench-Capon, editors, *Proceedings of the First International Conference on Computational Models of Argument*. IOS Press, 2006.

[11] T.J.M. Bench-Capon. Persuasion in practical argument using value based argumentation frameworks. *Journal of Logic and Computation*, 13(3):429–448, 2003.

[12] P. Besnard and A. Hunter. Practical First-order Argumentation. *Proc. of the 20th American National Conference on Artificial Intelligence (AAAI'2005)*, pages 590–595, 2005.

[13] P. Besnard and A. Hunter. *Elements of Argumentation*. MIT Press, 2008.

[14] C. Boutilier. Toward a logic for qualitative decision theory. In Jon Doyle, Erik Sandewall, and Pietro Torasso, editors, *Principles of Knowledge Representation and Reasoning*, pages 75–86, 1994.

[15] B.Phillips, C. Ball, D. Sackett, D. Badenoch, S. Straus, B. Haynes, and M. Dawes. Oxford centre for evidence-based medicine levels of evidence. http://www.cebm.net/index.aspx?o=1025.

[16] M. Bratman. *Intention, Plans, and Practical Reason*. CSLI, 1999.

[17] B. Buchanan and R. Smith. Fundamentals of Expert Systems. *Annual Review Computer Science*, 3:23–58, 1988.

[18] N. Buckley, I. Gibson, and H. Schipper. Cancer 2025: Optimizing the delivery of cancer care. *Expert Rev of Anticancer Therapy*, 4:S59–S62, 2004.

[19] Martin Caminada and Leila Amgoud. On the evaluation of argumentation formalisms. *Artif. Intell.*, 171(5-6):286–310, 2007.

[20] D. Cartwright and K. Atkinson. Political engagement through tools for argumentation. In S. Doutre P. Besnard and A. Hunter, editors, *Computational Models of Argument, Proceedings of COMMA 2008*, volume 172. of *Frontiers in Artificial Intelligence and Applications*, pages 116–127. IOS Press, Amsterdam, The Netherlands., 2008.

[21] A. Collins and M. Quillian. Retrieval Time From Semantic Memory. *Journal of Verbal Learning and Verbal Behavior*, 8:240–247, 1969.

[22] M. Cristofanilli, D. Hayes, G. Budd, M. Ellis, A. Stopeck, J. Reuben, G. Doyle, J. Matera, W. Allard, M. Miller, H. Fritsche, G. Hortobagyi, and L. Terstappen. Circulating tumor cells: A novel prognostic factor for newly diagnosed metastatic breast cancer. *J Clin Oncol*, 23:1420–1430, 2005.

[23] R. Davis, H. Shrobe, and P. Szolovits. What is a knowledge representation? *AI Magazine*, 14(1):17–33, 1993.

[24] Randall Davis. A dss for diagnosis and therapy. *SIGMIS Database*, 8(3):58–72, 1977.

[25] P. Dawid and J. Mortera. Coherent analysis of forensic identification evidence. *J. Roy. Statist. Soc. B*, 58:425–443., 1996.

[26] F. de Dombal, D. Leaper, J. Staniland, A. McCann, and J. Horrocks. Computer-aided diagnosis of acute abdominal pain. *BMJ*, 2:9–13, 1972.

[27] P. Dung. On the Acceptability of Arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77:321–357, 1995.

[28] Early Breast Cancer Trialists' Collaborative Group (EBCTCG). Effects of chemotherapy and hormonal therapy for early breast cancer on recurrence and 15-year survival: an overview of the randomised trials. *Lancet*, 14 May 2005.

[29] M. Eccles, E. McColl, N. Steen, N. Rousseau, J. Grimshaw, D. Parkin, and I. Purves. Effect of computerised evidence based guidelines on management of asthma and angina in adults in primary care: cluster randomised controlled trial. *BMJ*, 325:941 – 948, 2002.

[30] E.Emerson and E. Clarke. Characterizing correctness properties of parallel programs using fixpoints. In *Proceedings of the 7th Colloquium on Automata, Languages and Programming*, pages 169–181, London, UK, 1980. Springer-Verlag.

[31] J. Fox and S. Das. *Safe and Sound: Artificial Intelligence in Hazardous Applications*. MIT Press, Cambridge, MA/London, UK, 2000.

[32] A. Garcia and G. Simari. Defeasible logic programming: an argumentative approach. *Theory and Practice of Logic Programming*, 4(2):95 – 138, January 2004.

[33] M. Genesereth and N. Nilsson. *Logical Foundations of Artificial Intelligence.* Morgan Kaufmann, 1987.

[34] T. Gordon and N. Karacapilidis. The zeno argumentation framework. In *ICAIL '97: Proceedings of the 6th international conference on Artificial intelligence and law*, pages 10–18, New York, NY, USA, 1997. ACM Press.

[35] T. Gruber. Toward principles for the design of ontologies used for knowledge sharing. *International Journal of HumanComputer Studies*, 43:907–928, 1993.

[36] I. Horrocks, O. Kutz, and U. Sattler. The even more irresistible $\mathcal{SROIQ}$. In *Proc. of the 10th Int. Conf. on Principles of Knowledge Representation and Reasoning (KR 2006)*, pages 57–67. AAAI Press, 2006.

[37] I. Horrocks and P. Patel-Schneider. Reducing OWL entailment to description logic satisfiability. In *Proc. of the 2nd International Semantic Web Conference (ISWC)*, pages 17 – 29, 2003.

[38] I. Horrocks and P. Patel-Schneider. Reducing OWL Entailment to Description Logic Satisfiabilty. *Journal of Web Semantics*, 1(4):345 – 347, 31 October 2004.

[39] I. Horrocks, P. Patel-Schneider, S. Bechhofer, and D. Tsarkov. OWL rules: A proposal and prototype implementation. *Web Semantics: Science, Services and Agents on the World Wide Web*, 3(1):23–40, July 2005.

[40] Ian Horrocks, Peter F. Patel-Schneider, Harold Boley, Said Tabet, Benjamin Grosof, and Mike Dean. Swrl: A semantic web rule language combining owl and ruleml. http://www.w3.org/Submission/SWRL/, May 2004.

[41] D. Hunt, R. Haynes, S. Hanna, and K. Smith. Effects of computer-based clinical decision support systems on physician performance and patient outcomes: a systematic review. *Journal of the American Medical Association*, 280:1339–1346, October 1998.

[42] A. Hunter. Real arguments are approximate arguments,. In *Proceedings of the 22nd AAAI Conference on Artificial Intelligence*, pages 66–71. MIT Press., 2007.

[43] F. Ingrand, M. Georgeff, and A Rao. An architecture for real-time reasoning and system control. *IEEE Expert*, 7(2):33 – 44, 1992.

[44] Emery J, Walton R, Murphy M, Austoker J, Yudkin P, Chapman C, Coulson A, Glasspool D, and Fox J. Computer support for interpreting family histories of breast and ovarian cancer in primary care: comparative study with simulated cases. *BMJ*, 321:28–32, 2000.

[45] M. Johnston, K. Langton, R. Haynes, and A. Mathieu. Effects of computer-based clinical decision support systems on clinician performance and patient outcome. a critical appraisal of research. *Ann Intern Med.*, 120(2):135–42, 1994.

[46] Nikos Karacapilidis and Dimitris Papadias. Hermes: Supporting argumentative discourse in multi-agent decision making, 1998.

[47] J. Keynes. *General Theory of Employment, Interest and Money*. Macmillan, London, 1936.

[48] Holger Knublauch, Mark A. Musen, and Alan L. Rector. Editing description logic ontologies with the protégé owl plugin. In *In International Workshop on Description Logics*, 2004.

[49] L. Kohn, J. Corrigan, and M. Donaldson, editors. *To Err is Human*. IoM, 2000.

[50] P. Krause, S. Ambler, M. Elvang-Goranssan, and J. Fox. A logic of argumentation for reasoning under uncertainty. *Computational Intelligence*, 11:113–131, 1995.

[51] V. Lifschitz. *Principles of Knowledge Representation, Studies in Logic, Language and Information*, chapter Foundations of Declarative Logic Programming. CSLI Publishers, 1991.

[52] R. E. Fikes M. R. Genesereth. Knowledge interchange format, version 3.0 reference manual. Technical report, Computer Science Department, Stanford University, June 1992.

[53] John McCarthy. *Programs with Common Sense*. Her Majesty's Stationary Office, London, 1959.

[54] K. McPherson, C. Steel, and J Dixon. Breast cancer: Epidemiology. Risk factors and Genetics. *BMJ*, 321:624–628, 2000.

[55] Sanjay Modgil and Michael Luck. Argumentation based resolution of conflicts between desires and normative goals. In *ArgMAS*, pages 19–36, 2008.

[56] B. Motik, U. Sattler, and R. Studer. Query answering for OWL-DL with rules. *Journal of Web Semantics*, 3(1):41–61, 2005.

[57] Boris Motik, Rob Shearer, and Ian Horrocks. Optimized Reasoning in Description Logics using Hypertableaux, July 17–20 2007.

[58] P. Patel-Schneider, P. Hayes, and I. Horrocks. OWL Web Ontology Language Semantics and Abstract Syntax. Technical report, http://www.w3.org/TR/owl-semantics/, 10 February 2004.

[59] V Patkar, C Hurt, R Steele, S Love, A Purushotham, M Williams, R Thomson, and J Fox. Evidence-based guidelines and decision support services: a discussion and evaluation in triple assessment of suspected breast cancer. *British Journal of Cancer*, 95:1490–1496, 2006.

[60] J. Pollock. Rational cognition in OSCAR. In *ATAL '99: 6th International Workshop on Intelligent Agents VI, Agent Theories, Architectures, and Languages (ATAL),*, pages 71–90. Springer-Verlag, London, UK, 2000.

[61] D. Poole. Logical argumentation, abduction, and Bayesian decision theory: A Bayesian approach to logical arguments and its application to legal evidential reasoning. *Cardozo Law Review*, 22:1733–1745, 2002.

[62] H. Prakken and G. Sartor. Argument-based extended logic programming with defeasible priorities. In Pierre-Yves Schobbens, editor, *Working Notes of 3rd ModelAge Workshop: Formal Models of Agents*, Sesimbra, Portugal, 1996.

[63] M. Quinn and E. Allen. Changes in incidence of and mortality from breast cancer in england and wales since introduction of screening. *BMJ*, 311:1391–1395, 1995.

[64] R. Brachman and H. Levesque. *Knowledge Representation and Reasoning*. Morgan Kaufman,, 2004.

[65] A. Rector. A personal view of where i have been and where we might be going. *Artificial Intelligence in Medicine*, 23:111–127, 2001.

[66] C. Reed. Representing dialogic argumentation. *Knowledge Based Systems*, 19:22–31, 2006.

[67] M. Richards, I. Smith, and J. Dixon. Role of systemic treatment for primary operable breast cancer. *BMJ*, 309:1263–1366, 1994.

[68] D. Sackett, R. Haynes, P. Tugwell, and G. Guyatt. *Clinical Epidemiology*. Lippincott Williams and Wilkins, 2nd edition, 1991.

[69] D. Schum. *Evidential Foundations of Probabilistic Reasoning*. John Wiley & Sons, Toronto, 1994.

[70] E. Sirin, B. Parsia, B. Cuenca Graua, A. Kalyanpura, and Y. Katza. Pellet: A practical OWL-DL reasoner. *Journal of Web Semantics*, In Press.

[71] J. Sowa. *Knowledge representation: Logical, Philosophical and Computational Foundations*. Brooks/Cole California, 2000.

[72] D. Sutton and J. Fox. The syntax and semantics of PROforma. *J Am Med Inform Assoc.*, 10(5):433–43, 2003.

[73] P. Tolchinsky, K. Atkinson, P. McBurney, S. Modgil, and U. Cortés. Agents deliberating over action proposals using the proclaim model. In *Proceedings of the Fifth International Central and Eastern European Conference on Multi-Agent Systems (CEEMAS 2007)*, MultiAgent Systems and Applications V: Lecture Notes in Artificial Intelligence (LNAI) 4696, pages 32–41, 2007.

[74] L. van't Veer, S. Paik, and D. Hayes. Gene expression profiling of breast cancer: a new tumor marker. *J Clin Oncol*, 23:1631–1635, 2005.

[75] B. Verheij. Argue! - an implemented system for computer-mediated defeasible argumentation. In Han La Poutre and Jaap van den Herik, editors, *NAIC '98. Proceedings of the Tenth Netherlands/Belgium Conference on Artificial Intelligence*, pages 57–66. CWI, Amsterdam, 1998.

[76] R Walton, C Gierl, P Yudkin, H Mistry, M Vessey, and J Fox. Evaluation of computer support for prescribing (capsule) using simulated cases. *BMJ*, 315:791–795, 1997.

[77] J.H. Wigmore. *The Principles of Judicial Proof.* Little, Brown & Co., 1931.

[78] M. Wooldridge. *Reasoning about Rational Agents,.* MIT Press, Cambridge,MA, USA., 2000.

[79] Jeremy Wyatt and David Spiegelhalter. Evaluating medical expert systems: what to test and how? *Journal of Medical Informatics*, 1991.