

Harnessing Ontologies for Argument-based Decision-Making in Breast Cancer

Matt Williams*
London Research Institute,
Cancer Research UK
44 Lincoln Inn Fields, London WC2A 3PX
Matthew.Williams@cancer.org.uk

Anthony Hunter
Dept. Computer Science
University College London
Gower St., London, WC1E 6BT
A.Hunter@ucl.ac.uk

Abstract

We introduce a novel Ontology-based Argumentation Framework (OAF) that links a logic-based argumentation formalism and description logic ontologies. We show how these two formalisms can be tightly coupled by observing a few simple restrictions, and provides features not available in either formalism alone. Our work is evaluated in a large case study on decision-making in treatment choice in breast cancer, where rules are developed from the results of published clinical trials, and we present a small subset of this to demonstrate the use of the system. We show that OAF provides five advantages: (1) facilitating the clear use of shared definitions between multiple authors; (2) enabling us to match terms in the ontology and rules with those in the specific domain literature; (3) providing a close fit between structure of clinical trials and the structure of our rules; (4) delivering significant economies in the size of the rule-base compared to existing approaches; (5) allowing us take advantage of developments in both ontological and argumentative approaches. We also demonstrate that even a restricted language such as ours is sufficient to capture enough information to generate arguments that are useful for clinical practice. An early prototype implementation is available.

1. Introduction

The increase in medical knowledge has not been matched by an increase in our ability to use it, and evidence of the harm caused by medical care (98 000 deaths/year in the USA [16]) continues to mount. Because of this, Computer-based Decision Support Systems (CDSS) continue to appear attractive [15, 14]. Many prototype CDSS have concentrated on refining techniques for defeasible reasoning, but have not considered the structure of the domain

on which they operate. At the same time, there has been a drive to better characterise such domains, first in databases and now ontologies, whose reasoning mechanisms only allow us to draw strict conclusions. What is lacking is a technique to both (defeasibly) reason about a possible belief or action and to (strictly) reason about how these beliefs or actions relate to one another in a model of the world. Such requirements are necessary to be able to represent defeasible knowledge in a complex, real-world domain such as clinical medicine. For example, we need to defeasibly reason about the effects of a group of treatments, while at the same time knowing which treatments are in that group. Here we present a formalism that harnesses two existing knowledge representation (KR) approaches, argumentation and ontologies, to provide a formalism that has several attractive properties. Furthermore, our knowledge comes from the results of published clinical trials, and so the predicates in the ontology correspond to terms in the medical literature, while arguments are generated by drawing defeasible inferences based on the outcomes of these trials.

Argumentation offers an approach to reasoning with inconsistent knowledge (review in [5]), where an argument is a pair (support, claim), and a claim is a formula in a logical language and the support is a proof for the claim. For any given claim, inconsistencies in the knowledge base lead to arguments that are ‘pro’ and ‘con’ the claim. These arguments are said to attack one another and resolution is either through the development of different extensions, or the defeat of one argument by another. There are several different approaches to argumentation [17, 4, 7, 3], and our work extends the approach of Defeasible Extended Logic Programming (DeLP) [10]. This is based on extended logic programming, and uses rules to construct arguments; an implemented prototype is available for testing. OAF extends DeLP by ensuring that the predicates that are used in the rules are in the ontology, and that there is no predicate in the rules that does not appear in the ontology. We also redefine argument construction and conflict in terms of ontological entailment and unsatisfiability of formulae.

*This Author is supported as a Clinical Research Fellow by Cancer Research UK

An *ontology* is a model of the world. We use a description logic (DL) as the basis for our ontology, and primary reasoning tasks include deduction of subsumption of one concept (the subclass) by another (the superclass), membership of a set (class membership of individuals) and the satisfiability of formulae. There has been a great deal of work in recent years in the development of ontologies, and several extensive ontologies now exist in different biomedical domains, such as genetics (the Gene Ontology, > 50 000 terms), clinical medicine (SNOMED-CT, >350 000 terms) and cellular signalling (BioPAX).

Description Logics (DLs) are a subset of First-Order Logic [2], restricted to monadic and dyadic predicates (classes and properties) and with varying restrictions on the combination of different predicates. There are well described complexity results for several different description logics, and they provide the formal basis for the OWL family of Semantic Web ontology languages[19]. Such logics have two types of constants (*individuals* - instances of things in the world, and *datatypes* - dates, numbers, etc.) and three types of predicate (*classes* - sets of individuals, *object properties* - binary relations on individuals and *datatype properties* - binary relation on an individual and a datatype). For many theoretical and practical purposes it is useful to divide a DL ontology into non-ground formulae (the T-Box) and ground formulae (the A-Box). There are also well-developed tools for authoring and reasoning with such logics, such as Pellet [8] and Protege-OWL. However, Argumentation cannot be done in DL alone, as there is no equivalent in DL to the argumentative ideas of argument comparison and conflict.

Breast Cancer is the commonest non-skin cancer in women in the western world. Since many breast cancers are ‘driven’ by circulating sex hormones such as oestrogen, one of the treatments is the use of hormone receptor antagonists. Two notable ones are Tamoxifen and Anastrozole, and there have been extensive trials of different treatment courses of these drugs in patients with breast cancer. The effect of these drugs is greater in those women who have oestrogen receptor (ER) positive disease, as this is the target for the drugs, although the two drugs act on the receptor in very different ways.

Our approach is to use information and inferences from both the ontology and argumentation formalisms without combining them in a single knowledge-base. This is based upon four considerations: (1) The arguments are developed from rules drawn from the clinical literature; the ontology contains the same terms as the literature and rules, but the relationship between them is given by a combination of background knowledge and medical reference works, such as disease classification schemes [20]. (2) The predicates that can appear in each defeasible rule are defined in the ontology. This constraint provides a precise meaning for the

rules and hence the arguments. (3) The argumentation system constructs arguments from rules with antecedents that are satisfied when fielded as queries to the ontology. (4) Ontological reasoning is computationally more efficient than argumentative reasoning, and we maximise the use of more efficient DL reasoners by using them for the ontological information and restricting the argumentation system to only handling defeasible rules.

2. Formal Development

We start by defining an ontology, and ontological entailment. Throughout this section we shall use a running example about using hormone-receptor antagonists (e.g. Tamoxifen and Anastrozole) in breast cancer.

Definition 2.1. An *ontology* Δ is a set of well-formed formulae in a description logic, \mathcal{K} . The T-box of Δ is denoted Δ_T and the A-Box Δ_A . The symbol \vdash_{Ont} is used to denote ontological entailment in \mathcal{K} as defined in [2], with reference to the T-Box of the ontology Δ that we are using. An ontology is inconsistent, denoted $\vdash_{Ont} \top \subseteq \perp$, if for all interpretations of Δ , the interpretation is an empty model.

Example 2.1. For this paper, our ontology, Δ_{BC} , is based on the Breast Cancer domain¹. Although OAF is agnostic about the choice of DL, we have used OWL1.1 (corresponding to the DL *SRONTQ*) for this ontology. It consists of 89 classes (monadic predicates with leading capitals) that include Tamoxifen and Anastrozole (subclasses of DrugRegimes), where there are two subclasses of the former, Tamoxifen5Yr and Tamoxifen2Yr, as well as classes of Men and Women, both of which are People, and Diseases, such as Breast and Endometrial cancer, Osteoporosis and Fracture. We also need to be able to talk about Outcomes, and so we have classes of Increased and Decreased Overall Survival (IncreasedOS and DecreasedOS) as well as IncreasedRisk and ReducedRisk of Diseases, and various other personal attributes, such as MenopausalStatus. The ontology also contains 15 properties (binary predicates), either object properties to link individuals in these classes - thus hasDisease links a Person with a Disease, - and datatype properties that link individuals to datatypes - thus hasAge links an Outcome to an integer (where datatypes are defined by the XML datatype schema).

We now define the language of OAF in relation to some DL ontology. An OAF is a set of ground formulae and rules, but all the predicates that appear in OAF are found in the ontology, Δ .

Definition 2.2. For some ontology Δ in some DL language, \mathcal{K} , let $\mathcal{L}_{\mathcal{K}}$ be a logical language such that: C is the set of class names in \mathcal{K} ; R is the set of object property names in

¹Available at <http://acl.icnet.uk/~mw/BrCaReduced.owl>

\mathcal{K} ; S is the set of datatype property names in \mathcal{K} ; I is the set of named individuals in \mathcal{K} , denoted a, b, c ; D is the set of datatype value names in \mathcal{K} denoted $d_1 \dots d_n$; V is a set of variables denoted p, q, \dots, z , possibly with subscripts; $\#$ be the set of operators $\{\neg, \wedge\}$

Example 2.2. So for $\mathcal{L}_{\Delta_{BC}}$, C includes the monadic predicates Treatments, People, Tamoxifen, etc.

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

Definition 2.3. An *atom* is any of the following: $p(t)$ where $p \in C$ and $t \in I \cup V$; $p(t_1, t_2)$ where $p \in R$ and $t_1, t_2 \in V \cup I$; $p(t_1, t_2)$ where $p \in S$ and $t_1 \in V \cup I, t_2 \in D$. The set of all atoms is denoted \mathcal{A} ; a *literal* is an atom or a negated atom (i.e. if $p(t)$ is a monadic atom, then $\neg p(t)$ is a literal); a *formula* is a literal or conjunction of literals. The set of all formulae is denoted \mathcal{F} , and the set of all ground formulae (i.e. those that do not contain any variables) is denoted \mathcal{F}_G

Definition 2.4. A *defeasible rule* is of the form: $L: \beta \Rightarrow \alpha$ where $\alpha, \beta \in \mathcal{F}_G$ and L is an alphanumeric label. We refer to α as the *Body* of the rule, β as the *Head* of the rule and L as the *Label*. The functions $Head(r) = \alpha, Body(r) = \beta$ return parts of some rule, r .

As stated above, our rules are ground. However, following DeLP [10] we will sometimes use “schematic rules” with variables. Given a “Schematic Rule” R , $Ground(R)$ stands for all the ground instances of R . Variables in schematic rules are denoted by the letters x, y, z .

Example 2.3. Consider the treatment of breast cancer with Hormone Receptor antagonists. There are four major drugs in this class, but we shall concern ourselves with only anastrozole and tamoxifen. There have been trials with either two years or five years of tamoxifen and in our example below, one rule concerns the use of 5 years of tamoxifen, one the use of tamoxifen (i.e. either 2 years or 5 years) and one anastrozole:

r_a : People(x)
 \wedge hasDisease(x,y) \wedge BreastCancer(y)
 \wedge Tamoxifen5Yr(z)
 \Rightarrow hasTreatment(x,z)

r_b : People(x)
 \wedge hasDisease(x,y) \wedge BreastCancer(y)
 \wedge Tamoxifen(z)
 \Rightarrow hasTreatment(x,z)

r_c : People(x)
 \wedge hasDisease(x,y) \wedge BreastCancer(y)
 \wedge Anastrozole(z)
 \Rightarrow hasTreatment(x,z)

Definition 2.5. An *Ontology-based Argumentation Framework* (OAF), Ω_Δ , for some ontology $\Delta \subset \mathcal{K}$, is a finite set of defeasible rules (‘rules’), Θ and formulae in \mathcal{F} for $\mathcal{F} \subset \mathcal{L}_\mathcal{K}$. We denote the set of ground formulae (‘facts’) in Ω_Δ as Γ , such that $\Gamma \subset \mathcal{F}_G$.² We denote a simple Ω_Δ as $(\Theta \cup \Gamma)$, but shall often write $\Omega_\Delta = (\Theta, \Gamma)$.

An OAF therefore consists of a set of defeasible rules and facts, where the facts are ground formulae in the A-Box of ontology. As we develop the running example, we shall consistently refer to the same ontology, Δ_{BC} , although we may vary the rules and facts to demonstrate different aspects of the work.

2.1. Ontological Reasoning

The coupling of the ontology and argumentation in OAF allows us to argue without worrying about the details of the ontological formalism. We access information in the ontology through a set of two relations, $\text{conflict}(\phi, \psi)$ and $\text{entails}(\phi, \psi)$, and in order to do so, we must translate a formulae in OAF into an ontological query. We do this using the ‘Rolling Up’ procedure for conjunctive A-Box queries described by Horrocks [13], which we denote ρ . This translation maps arbitrary conjunctions of formulae in OAF to existentially quantified formulae in DL, where monadic predicates are classes and binary predicates are properties, as the next example illustrates:

Example 2.4. Consider the OAF formula: $\phi = \text{People}(x) \wedge \text{hasDisease}(x,y) \wedge \text{BreastCancer}(y)$; then $\rho(\phi) = \text{People} \sqcap \exists \text{hasDisease.BreastCancer}$.

Note the typically terse DL syntax; translation of such formulae into conventional first-order logic yields formulae with one free variable. For more details, see [2].

We harness the ontology for reasoning about literals arising in the rules; we are principally interested in whether or not two formulae are mutually inconsistent, or whether one entails the other, as given by the following definitions:

Definition 2.6. For two formulae, $\phi, \psi \in \mathcal{F}$ $\text{entails}(\phi, \psi)$ holds iff $\Delta_T \cup \rho(\phi) \vdash_{Ont} \rho(\psi)$.

Definition 2.7. For any two formulae in $\phi, \psi \in \mathcal{F}$ $\text{conflict}(\phi, \psi)$, holds iff $\Delta_T \cup (\rho(\phi) \sqcap \rho(\psi)) \vdash_{Ont} \perp$ and $\Delta_T \cup (\rho(\phi)) \not\vdash_{Ont} \top \subseteq \perp$ and $\Delta_T \cup (\rho(\psi)) \not\vdash_{Ont} \top \subseteq \perp$

In general, $\text{entails}(\phi, \psi)$ holds in two circumstances. The first is if ϕ is a subclass of ψ , and the second is if ϕ is equivalent to ψ . Also note that conflict is symmetric i.e. $\text{conflict}(\phi, \psi) \leftrightarrow \text{conflict}(\psi, \phi)$, whereas as entailment may not be.

² Γ therefore corresponds to the A-Box of the ontology

Example 2.5. Using predicates from the example above, $\text{entails}(\text{Tamoxifen5Yr}(x), \text{Tamoxifen}(x))$ and $\text{conflict}(\text{hasTreatment}(x,z) \wedge \text{Tamoxifen}(z), \text{hasTreatment}(x,z) \wedge \text{Anastrozole}(z))$ both hold for any grounding of x and z .

3. Arguments

So far we have introduced the language of our system and rules. However, our aim is to develop arguments about the literals in our language, and it is to this that we now turn. Intuitively, an argument is a set of inferences (the support) that allows one to infer a formula (the claim)³. This work is broadly similar to that in DeLP, on which our system is based, with the exception that definitions of entailment and conflict are defined with reference to the ontology, rather than via strict rules. Other definitions follow those in DeLP, but have been amended to take this into account.

Definition 3.1. Let $\Omega_\Delta = \{\Theta, \Gamma\}$ and ϕ be a ground formula in Γ . A *defeasible derivation* of ϕ from Ω_Δ , denoted $\Omega_\Delta \vdash_{def} \phi$, consists of a finite sequence ϕ_1, \dots, ϕ_n of ground formulae, and each formula is in the sequence because ϕ_i is in Γ or there exists a rule, $r \in \Theta$, whose head contains ϕ_i and body ϕ_1, \dots, ϕ_k and every literal of the body is either an element ϕ_j of the sequence appearing before ϕ_i ($j < i$) or every literal of the body is either entailed by the element ϕ_j of the sequence appearing before ϕ_i ($j < i$), such that $\text{entails}(\phi_j, \phi_i)$ holds, or is entailed by some ϕ_j in Γ .

Definition 3.2. A set of rules Θ is *contradictory* iff there exists a defeasible derivation for formulae ϕ_1 and ϕ_2 from (Θ, Γ) such that $\text{conflict}(\phi_1, \phi_2)$ holds.

Definition 3.3. Let $\phi \in \Gamma$, and $\Omega_\Delta = (\Theta, \Gamma)$. We say that $\text{Label} : \langle A, \phi \rangle$ is an *argument* for ϕ , if A is a set of defeasible rules from Θ , such that (1) there exists a defeasible derivation for ϕ from $\Gamma \cup A$, (2) the set $\Gamma \cup A$ is non-contradictory, and (3) A is minimal: there is no proper subset A' of A such that A' satisfies conditions (1) and (2).

In short, an argument $\text{Label} : \langle A, \phi \rangle$, is a minimal non-contradictory set of defeasible rules, obtained from a defeasible derivation for a given ground formula, ϕ . We also call ϕ the *claim* supported by A . We shall also frequently refer to the rules in the support by their labels rather than reproducing the entire rule.

Example 3.1. Continuing example 2.3 above, let $\Omega_\Delta = (\Theta, \Gamma)$, where Γ is the following set:

```
Woman(MsJones)
hasDisease(MsJones,D1)
ERPositiveDisease(D1)
```

³Some formalisms include both the premises and the inferences in the support - ours does not

```
EarlyBreastCancer(D1)
Tamoxifen5Yr(TamCourseA)
Tamoxifen2Yr(TamCourseB)
Anastrozole(AnastrozoleCourseA)
ZoledronicAcid(ZolendronateA)
```

Then for all groundings of x and z :

```
entails(Woman(x),People(x))
entails(EarlyBreastCancer(x),BreastCancer(x))
entails(Tamoxifen5Yr(x),Tamoxifen(x))
entails(Tamoxifen2Yr(x),Tamoxifen(x))
conflict(hasTreatment(x,z) \wedge Tamoxifen(z),
hasTreatment(x,z) \wedge Anastrozole(z))
```

hold, and we can form the following arguments:

```
A1: \langle \{r_a\}, hasTreatment(MsJones, TamCourseA) \rangle
A2: \langle \{r_b\}, hasTreatment(MsJones, TamCourseA) \rangle
A3: \langle \{r_b\}, hasTreatment(MsJones, TamCourseB) \rangle
A4: \langle \{r_c\}, hasTreatment(MsJones, AnastrozoleCourseA) \rangle
```

We shall use this Γ again in the next section.

Definition 3.4. An argument $\langle B, \psi \rangle$ is a sub-argument of $\langle A, \phi \rangle$ if $B \subseteq A$.

Definition 3.5. Let $\Omega_\Delta = (\Theta, \Gamma)$. We say that two ground formulae ϕ and ϕ_1 in Ω *disagree* iff $\text{conflict}(\phi, \phi_1)$ holds.

Definition 3.6. Let \succ 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 $\langle A_1, \phi_1 \rangle$ will be preferred over $\langle A_2, \phi_2 \rangle$ iff there exists at least one rule $r_a \in A_1$, and one rule $r_b \in A_2$, such that $r_a \succ r_b$ and there is no $r'_b \in A_2$ and $r'_a \in A_1$ such that $r'_b \succ r'_a$.

The derivation of preferences between rules can be a complex matter. In our domain, we consider factors such as the type of evidence, the number of patients in a trial, the quality of the journal, the geographical location of the patients and the cost of the intervention. For this paper we assume that it is given explicitly via a simple ordering on the elements of Θ .

Definition 3.7. Let $\langle A_1, \phi_1 \rangle$ and $\langle A_2, \phi_2 \rangle$ be two arguments. Then $\langle A_1, \phi_1 \rangle$ is a *proper defeater* for $\langle A_2, \phi_2 \rangle$ at ϕ iff there exists a sub-argument $\langle A, \phi \rangle$ of $\langle A_2, \phi_2 \rangle$ such that $\langle A_1, \phi_1 \rangle$ counter-argues $\langle A, \phi \rangle$ at ϕ and $\langle A_1, \phi_1 \rangle \succ \langle A, \phi \rangle$.

Of course, it may not be the case that the two arguments are related by a preference ordering

Definition 3.8. 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 $\langle A_2, \phi_2 \rangle$ such that $\langle A, \phi_1 \rangle$ counter-argues $\langle A_2, \phi_2 \rangle$ at ϕ and $\langle A_1, \phi_1 \rangle$ is unrelated to by a preference-ordering to $\langle A, \phi \rangle$ i.e. $\langle A_1, \phi_1 \rangle \not\prec \langle A, \phi \rangle$ and $\langle A, \phi \rangle \not\prec \langle A_1, \phi_1 \rangle$.

Definition 3.9. 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 3.2. Continuing our example above, we can see that A_1 and A_4 conflict, as do A_2 and A_3 . Given the preference ordering $r_a \succ r_b \succ r_c$, we can see that A_1 is a proper defeater for A_4 , but A_2 is a blocking defeater for A_3 (and vice versa).

Definition 3.10. Let $\langle A_0, \phi_0 \rangle$ be an argument in Ω , such that $\Omega \vdash_{def} \langle A_0, \phi_0 \rangle$. An **argumentation line** for $\langle A_0, \phi_0 \rangle$ is a sequence of arguments from Ω , denoted $\Lambda = [\langle A_0, \phi_0 \rangle, \langle A_1, \phi_1 \rangle, \dots]$ where each element in the sequence $\langle A_i, \phi_i \rangle$, is a defeater of its predecessor, $\langle A_{i-1}, \phi_{i-1} \rangle$.

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

Definition 3.11. Let $\Lambda = [\langle A_0, \phi_0 \rangle, \langle A_1, \phi_1 \rangle, \dots]$ be an argumentation line. We define the set of **supporting arguments** $\Lambda_s = \{\langle A_0, \phi_0 \rangle, \langle A_2, \phi_2 \rangle\}$ and the set of **interfering arguments**, $\Lambda_I = \{\langle A_1, \phi_1 \rangle, \langle A_3, \phi_3 \rangle\}$

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

Definition 3.12. Two arguments $\langle A_1, \phi_1 \rangle$ and $\langle A_2, \phi_2 \rangle$ are **concordant** if $\{A_1 \cup A_2\}$ are not contradictory.

The notion of concordant arguments lets us enforce a restriction that in general lines of argumentation must be consistent; we cannot have an argument that defends the root argument while at the same time attacking it. A similar intuition is captured in Dung's [7] definition of a 'controversial' argument.

Definition 3.13. Let $\Lambda = [\langle A_0, \phi_0 \rangle, \langle A_1, \phi_1 \rangle, \dots]$ be an argumentation line. Λ is an **acceptable argumentation line** iff Λ is a finite sequence and the set Λ_s of supporting arguments is concordant and the set Λ_I of interfering arguments is concordant and 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 and 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 $\langle A_i, \phi_i \rangle$.

Definition 3.14. Let $\langle A_0, \phi_0 \rangle$ be an argument from Ω_Δ . A **dialectical tree** for $\langle A_0, \phi_0 \rangle$, denoted $\mathcal{T}_{\langle A_0, \phi_0 \rangle}$ is defined as follows: The root of the tree is labeled with $\langle A_0, \phi_0 \rangle$; then let N be a non-root node of the tree labeled $\langle A_n, \phi_n \rangle$ and

$\Lambda = [\langle A_0, \phi_0 \rangle, \langle A_1, \phi_1 \rangle, \dots]$ the sequence of labels of the path from the root to N . Let $\langle B_0, \psi_0 \rangle, \langle B_1, \psi_1 \rangle, \dots, \langle B_k, \psi_k \rangle$ 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 $\Lambda' = [\langle A_0, \phi_0 \rangle, \langle A_1, \phi_1 \rangle, \dots, \langle A_n, \phi_n \rangle, \langle B_i, \psi_i \rangle]$ is acceptable, then the node N has a child N_i labeled $\langle B_i, \psi_i \rangle$. If there is no defeater for $\langle A_n, \phi_n \rangle$ 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 is not always a dialectical tree. Given an acceptable argumentation line where a defeater $\langle A, \phi \rangle$ will not be included because it would make the line unacceptable, there might be a subsequence of the line where $\langle A, \phi \rangle$ could be included.

Definition 3.15. 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: All leaves in $\mathcal{T}_{\langle A, \phi \rangle}$ are marked as 'U's (undefeated) in $\mathcal{T}_{\langle A, \phi \rangle}^*$ and 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' (Defeated). The node $\langle B, \psi \rangle$ will be marked as 'D' in $\mathcal{T}_{\langle A, \phi \rangle}^*$ iff it has at least a child marked as 'U'.

Example 3.3. For an example of a dialectical tree, please see Fig. 1.

Definition 3.16. 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'. We 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**.

Example 3.4. Continuing Example 3.1, if we construct the tree for A_3 , we know that it is attacked by A_1 , which is undefeated, and thus A_3 is defeated. Therefore the argument $A_3 : \langle \{r_b\}, \text{hasTreatment}(\text{MsJones}, \text{TamCourseB}) \rangle$ is unwarranted.

In the rest of this paper we shall briefly examine the effect of incorporating an ontology into the argumentation formalism. We shall do this via a case study and discussion.

4. Case Study

To investigate the use of OAF with real data, we have applied it to decision-making in early breast cancer. We have identified a set of 57 papers referenced by one section of the National Cancer Institute's (NCI, USA [1]) breast cancer guideline. These include randomised controlled trials RCTs, guidelines and meta-analyses of trials. For details on an approach to analysing such papers see [11]. From these we have developed 117 defeasible rules and an

ontology of 190 classes and 31 properties of which Δ_{BC} is a subset. Here we present a subset of the case study. Where the paper is a clinical trial of a treatment which has a beneficial effect, we have encoded that as a rule for the treatment; where there is a harmful effect, we have used a rule against the treatment. It is our hypothesis that we should be able to reconstruct the advice given in the guideline by reasoning directly with the evidence. The predicates used should be familiar from the previous examples, with the exception of ERPositive disease, which refers to the oestrogen receptor status of the tumour (see Sec. 1) and Zoledronic acid, which is a drug that increases bone density and reduces bone-based metastasis of the tumour. We have also developed an early prototype implementation in Java⁴.

Let Ω_{Δ} be the OAF where Γ is as for Example 3.1 and Θ contains exactly the following seven rules:

r_1 : (Women(x)
 \wedge hasDisease(x,y) \wedge EarlyBreastCancer(y)
 \wedge ERPositiveDisease(y)
 \wedge Tamoxifen5Yr(z)
 \Rightarrow hasTreatment(x,z))

r_2 : (Women(x)
 \wedge hasDisease(x,y) \wedge EarlyBreastCancer(y)
 \wedge ERPositiveDisease(y)
 \wedge Tamoxifen2Yr(z)
 \Rightarrow hasTreatment(x,z))

r_3 : (Women(x)
 \wedge hasDisease(x,y) \wedge BreastCancer(y)
 \wedge Tamoxifen(z)
 $\Rightarrow \neg$ hasTreatment(x,z))

r_4 : (Women(x)
 \wedge hasDisease(x,y) \wedge BreastCancer(y)
 \wedge Tamoxifen5Yr(z)
 $\Rightarrow \neg$ hasTreatment(x,z))

r_5 : (Women(x)
 \wedge hasDisease(x,y) \wedge BreastCancer(y)
 \wedge hasMenopausalState(x,PostMenopausal)
 \wedge Tamoxifen(z)
 \Rightarrow hasTreatment(x,z))

r_6 : (Women(x)
 \wedge hasDisease(x,y) \wedge BreastCancer(y)
 \wedge Anastrozole(z)
 $\Rightarrow \neg$ hasTreatment(x,z))

r_7 : (Women(x)
 \wedge hasDisease(x,y) \wedge BreastCancer(y)
 \wedge ZoledronicAcid(z)
 \Rightarrow hasTreatment(x,z))

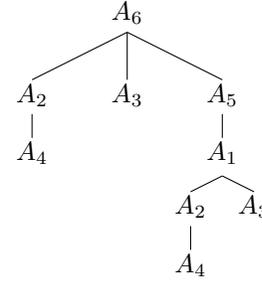


Figure 1. The Argument Tree for A_6

From Ω , we make the following arguments:

A_1 : $\{r_1\}$, hasTreatment(MsJones, TamCourseA)
 A_2 : $\{r_2\}$, hasTreatment(MsJones, TamCourseB)
 A_3 : $\{r_3\}$, \neg hasTreatment(MsJones, TamCourseA)
 A_4 : $\{r_3\}$, \neg hasTreatment(MsJones, TamCourseB)
 A_5 : $\{r_4\}$, \neg hasTreatment(MsJones, TamCourseA)
 A_6 : $\{r_5\}$, hasTreatment(MsJones, TamCourseA)
 A_7 : $\{r_5\}$, hasTreatment(MsJones, TamCourseB)
 A_8 : $\{r_6\}$, \neg hasTreatment(MsJones, AnastrozoleCourseA)
 A_9 : $\{r_7\}$, hasTreatment(MsJones, ZoledronateA)

For this example, we derived preferences between the rules based on the size of the study population described in the paper from which the rules are drawn. This gives us:

$$r_3 > r_2 > r_1 > r_4 > r_5 > r_6$$

We can then use this to develop our dialectical trees. As an example, we give the tree for A_6 , above. On marking the tree (as in 3.15), there is an undefeated leaf node that directly attacks the root node; as a result, the root node is defeated, and thus A_6 is unwarranted.

When compared to the breast cancer guideline from which the references were taken, the arguments that we construct are found in the guideline. However, not all of the arguments in the guideline are produced by our system, for three reasons. The first is that the guideline fails to reference certain inferences that it makes (e.g. from increased disease specific survival to increased overall survival). The second is that the guideline is more “relaxed” about the patient characteristics than the trials explicitly justify (e.g. a trial on *premenopausal* women with early stage breast cancer is used as the reference for a statement about *all* women with early stage breast cancer). The third is that where a treatment has mixed effects (as most do) our system often fails to correctly them. The first two results are interesting,

⁴Available for download at <http://acl.icnet.uk/~mw/swafi>

as they allow us to pinpoint exactly where the guideline fails to supply the evidence to support its statements, and thus we can use our approach as a way of checking such guidelines. The third is important as it highlights a weakness in current formalisations of argumentation.

5. Advantages of OAF

Our approach is based on an existing argumentation formalism, extended to incorporate an ontology. We argue that this has five benefits, summarised below.

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 `BreastCancer(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). In addition, since we allow only the predicates that appear in the ontology to appear in the rules, we make the process of authoring rules much easier; it becomes a matter of choosing the appropriate predicates from the ontology, and given the enforced simplicity of the language (restricted to monadic and dyadic predicates) we limit the choice.

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.

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 Trialists' Collaborative Group meta-analysis of trials of Tamoxifen [9] (the source of r_1 and r_2) included the table in Fig 2. Here we see a summary of the results of tri-

als 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 are can help one quickly understand the results of the paper.

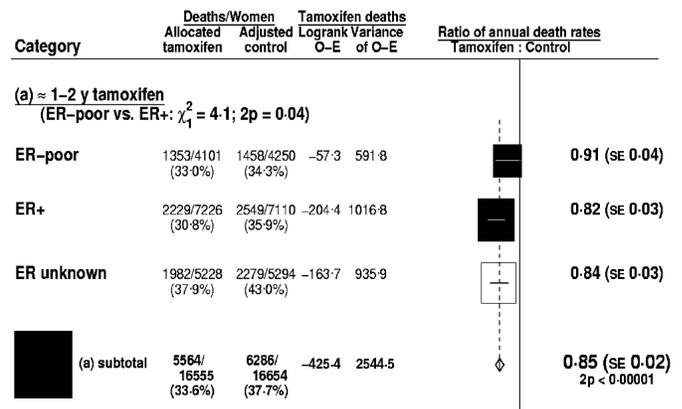


Figure 2. Forest Plot for Tamoxifen trials

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 [20].

Efficiency Our approach allows us to use the ontology to ease rule authoring. Consider r_1 :

```

Women(x)
^ hasDisease(x,y) ^ EarlyBreastCancer(y)
^ ERPositiveDisease(y)
^ Tamoxifen5Yr(z)
=>hasTreatment(x,z)

```

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. Such class-based subsumption allows us (for r_1) to write one less rule; More generally, however, if the n terms in a rule each subsume m others (where a leaf node subsumes

itself, $m=1$), then the reduction in rules is given by $m_1 \times m_2 \dots \times m_n$. This allows us to relate the structure of the ontology to the expected efficiency gains in the rulebase.

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 [12]. However, there are both practical and technical reasons for not doing so. Firstly, it 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 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 improvements in reasoning algorithms for argumentation.

6. Discussion

DL ontologies, and integration of ontologies with rules has received much attention in the literature. However, there has been much less on integrating ontologies and argumentation systems. The work that exists either generates a single knowledge-base [12] or uses a very simple form of argumentation without studying argument conflict and resolution [6]. Other methods of handling uncertainty in OWL (such as FuzzyOWL [18]) handle uncertainty without resolving conflicting information. We have introduced a novel argumentation framework, which allows us to incorporate ontological information into our arguments. We have evaluated this in a large case study, of which we have presented a small subset here. We have also shown how such an approach can enhance consistency, transparency and efficiency in a domain that seems well suited to our approach. It also allows us to take advantages of improvements in both fields, and an early prototype implementation is available. Our future work will focus on four main areas: 1) theoretical work on developing a formalisation of the practical syllogism and 2) abstracted arguments, such as for Tamoxifen (in general), and practical work on 3) developing the implementation to allow for caching of arguments and 4) provision of OAF as a server-based service.

References

[1] NCI Breast Cancer PDQ (Stage I, II, IIIA, and Operable IIIC Breast Cancer), 30 May 2007.
 [2] F. Baader, D. Calvanese, D. McGuinness, D. Nardi, and P. Patel-Schneider, editors. *The Description Logic Hand-*

book: Theory, Implementation and Applications. Cambridge University Press, Jan. 2003.
 [3] T. J. M. Bench-Capon. Persuasion in practical argument using value-based argumentation frameworks. *J. Log. Comput.*, 13(3):429–448, 2003.
 [4] P. Besnard and A. Hunter. A logic-based theory of deductive arguments. *Artificial Intelligence*, 128:203–235, 2001.
 [5] C. Chesnevar, A. Maguitman, and R. Loui. Logical models of argument. *ACM Comput. Surv.*, 32(4):337–383, 2000.
 [6] Y. S. C. Tempich, H. S. Pinto. An Argumentation Ontology for Distributed, Loosely-controlled and evolving Engineering processes of ontologies (DILIGENT). In *European Semantic Web Conference*, pages 241–256, 2005.
 [7] 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.
 [8] E. Sirin, B. Parsia, B. Cuenca Graua, A. Kalyanpura, and Y. Katza. Pellet: A practical OWL-DL reasoner. *Journal of Web Semantics*, In Press.
 [9] 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.
 [10] A. Garcia and G. Simari. Defeasible logic programming: an argumentative approach. *Theory and Practice of Logic Programming*, 4(2):95 – 138, Jan. 2004.
 [11] T. Greenhalgh. *How to Read a Paper: The Basics of Evidence Based Medicine*. BMJ Books, London, 2nd edition, Nov. 2000.
 [12] Grigoris Antoniou and Antonis Bikakis and Gerd 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.
 [13] I. Horrocks and S. Tessaris. A conjunctive query language for description logic ABoxes. In *Proc. of the 17th Nat. Conf. on Artificial Intelligence (AAAI 2000)*, pages 399–404, 2000.
 [14] 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. *JAMA*, 280:1339–1346, October 1998.
 [15] 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.
 [16] L. Kohn, J. Corrigan, and M. Donaldson, editors. *To Err is Human*. IoM, 2000.
 [17] P. Krause, S. Ambler, M. Elvang-Goransson, and J. Fox. A logic of argumentation for reasoning under uncertainty. *Computational Intelligence*, 11:113–131, 1995.
 [18] C. L. M. Gao. Extending OWL by fuzzy description logic. In *Tools with Artificial Intelligence, 2005.*, 14 Nov. 2005.
 [19] D. McGuinness and F. van Harmelen. OWL Web Ontology Language Overview. <http://www.w3.org/TR/owl-features/>, 10 Feb. 2004.
 [20] L. Sobin and C. Wittekind, editors. *TNM Classification of Malignant Tumours, 6th edition*. John Wiley & Sons, Hoboken, New Jersey, 2002.