Persuasion and Value in Legal Argument

TREVOR BENCH-CAPON, KATIE ATKINSON and ALISON CHORLEY,
Department of Computer Science, University of Liverpool, Liverpool L69 3BX, UK. Email: {tbc,katie,alison}@csc.liv.ac.uk

Abstract

In this paper we consider legal reasoning as a species of practical reasoning. As such it is important both that arguments are considered in the context of competing, attacking and supporting arguments, and that the possibility of rational disagreement is accommodated. We present two formal frameworks for considering systems of arguments: the standard framework of Dung, and an extension which relates arguments to values allowing for rational disagreement. We apply these frameworks to modelling a body of case law, explain how the frameworks can be generated to reconstruct legal reasoning in particular cases, and describe some tools to support the extraction of the value related knowledge required from a set of precedent cases.

Keywords: Argumentation, legal reasoning, practical reasoning, argumentation frameworks, argumentation schemes, critical questions, legal theory construction.

1 Introduction

Legal reasoning is typically directed towards the resolution of some disagreement. Sometimes it may be a matter of fact that is in dispute: such matters are resolved by presenting evidence to a jury. But when the dispute turns on a point of law, there is no fact of the matter: the court must choose which arguments they will follow, which position they will adopt. In this situation persuasion, not demonstration or proof, is the central notion: counsels for the parties attempt to persuade the court, and the judge writing the decision attempts to persuade any superior court to which the decision may be appealed, and, ultimately, public opinion. Because we are dealing with persuasion about a choice to be made, we cannot expect to be able to resolve the matter through a logically coercive argument:

As Perelman, the originator of the *New Rhetoric*, puts it:

> If men oppose each other *concerning a decision to be taken*, it is not because they commit some error of logic or calculation. They discuss apropos the applicable rule, the ends to be considered, the meaning to be given to values, the interpretation and characterization of facts. [24, p. 150] (*italics ours*).

A similar point is made by the philosopher John Searle:

> Assume universally valid and accepted standards of rationality. Assume perfectly rational agents operating with perfect information, and you will find that rational disagreement will still occur; because, for example, the rational agents are likely to have different and inconsistent values and interests, each of which may be rationally acceptable. [25, p. xv]

Both Perelman and Searle recognize that there may be complete agreement on facts, logic, which arguments are valid, which arguments attack one another and the rules of fair debate, and yet still disagreement as to the correct decision. We can illustrate such disagreements with a couple of simple examples.

Vol. 15 No. 6,
doi:10.1093/logcom/exi058

One choice that any government must make is to decide on an appropriate rate of income tax. Typically there will be an argument in favour of increasing the rate of taxation, since this progressive form of taxation will reduce income inequalities. Against this, it can be argued that a decrease in taxation will promote more enterprise, increasing Gross National Product, and so raising the absolute incomes of everyone. It is possible to see both these arguments as valid, since both supply a reason to act: and yet a choice must be made, since the actions are incompatible. Which choice is made will depend on whether the chooser prefers equality or enterprise in the particular circumstances with which he is confronted. Two parties may be in agreement as to the consequences of a movement in the tax rate, and yet disagree as to the choice to be made because they differ in their fundamental aspirations. Different people will prize social values differently, and one may prefer equality to enterprise, while another prefers enterprise to equality. Thus while both arguments are agreed to be valid, one *audience* will ascribe more force to one of the arguments, while a different *audience* will make a different choice. In such cases these different audiences will rationally disagree, and agreement can only be reached by coming up with additional arguments which convince all audiences *in terms of their own preferences*, or by converting those which disagree to a different appraisal of social values.

A similar situation arises in law if we consider the question of how wide should be the discretion allowed to judges. On the one hand, consistency is promoted if the law is clear and the letter of the law is applied. This, however, will inevitably mean that in some cases such application of the law will give rise to decisions which seem hard, or even unjust, given the facts of a particular case. So some would argue that justice would be served by giving the judge considerable freedom to apply the spirit rather than the letter of the law. Two competing notions of what is fair are here: that like cases should be treated alike, as against the need to consider the particular facts of individual cases, and a recognition that one size does not fit all. Again rational people can incline to either argument, while recognizing that both have merit.

To explain and model such disagreements we need to go beyond standard logic, since we must allow competing arguments to be accepted and yet choose one of them to govern the decision: as Perelman argues in the preface to *New Rhetoric* [24, p. 10]:

> Logic underwent a brilliant development during the last century when, abandoning the old formulas, it set out to analyze the methods of proof used effectively by mathematicians. ... One result of this development is to limit its domain, since everything ignored by mathematicians is foreign to it. Logicians owe it to themselves to complete the theory of demonstration obtained in this way by a theory of argumentation.

In this spirit we will discuss in this paper how argumentation can be used to model such disagreements and provide techniques for persuasion in such contexts. Our starting point will be two formal models of argumentation. That of Dung [22], models argumentation at its most abstract, and therefore does not accommodate any notion of value. The Value Based Argumentation Framework (VAF) of Bench-Capon [10] develops Dung's framework so that arguments can be related to underlying values and disagreements between audiences resulting from different preferences among values can be modelled. Section 3 will show how these frameworks can be applied to a body of case law. Section 4 discusses how we can generate the arguments needed to instantiate a VAF. Section 5 will explore deeper still and suggest how theories of a case law domain connecting facts and values and which supply the material to generate such arguments can be derived. Additionally, we describe some tools designed to support this process. We will conclude the paper with a discussion of some interesting issues that arise in connection with this work.

2 Formal models of argumentation

In this section we describe the Argument Framework (AF) of Dung [22] and the Value Based Argumentation Framework (VAF) of Bench-Capon [10].

For Dung the notion of an argument is as abstract as it can be: arguments are characterized only by the arguments they attack and are attacked by. This is especially suitable for modelling informal, natural language arguments, since the arguments are unconstrained in form, and there are no restrictions on what we can choose to count as an attack of one argument on another.

A formal definition of an Argumentation Framework, and the central notions concerning Argumentation Frameworks, is given as Definition 2.1.

DEFINITION 2.1
An Argumentation Framework (*AF*) is a pair $AF = \langle AR, A \rangle$, where *AR* is a set of arguments and $A \subset AR \times AR$ is the attack relationship for *AF*. A comprises a set of ordered pairs of distinct arguments in *AR*. A pair $\langle x,y \rangle$ is referred to as "*x attacks y*".

For *R*, *S*, subsets of *AR*, we say that

(a) $s \in S$ is attacked by *R* if there is some $r \in R$ such that $\langle r,s \rangle \in A$.

(b) $x \in AR$ is *acceptable* with respect to *S* if for every $y \in AR$ that attacks *x*, there is some $z \in S$ that attacks *y* (i.e. *z*, and hence *S*, defends *x* against *y*).

(c) *S* is *conflict free* if no argument in *S* is attacked by any other argument in *S*.

(d) A conflict free set is *admissible* if every argument in *S* is acceptable with respect to *S*.

(e) *S* is a *preferred extension* if it is a maximal (with respect to set inclusion) admissible subset of *AR*.

Note that in (b), *z* may be attacked by some argument, either a member of S, or outside of S. Thus while by (b) an argument may be acceptable with respect to S, its overall acceptability requires S to be admissible: then by (c) *z* will not be attacked by a member of S, and by (d) if *z* is attacked by an argument outside S it will have a defender within S. Note that an argument may defend itself. The key notions are the *admissible set*, which represents a group of arguments free from internal conflict and able to defend themselves against external attacks, and the *preferred extension* which represents a position which is:

- internally consistent
- can defend itself against all attacks
- cannot be further extended without becoming inconsistent or open to attack.

From Dung [22] we know that every *AF* has a preferred extension (possibly the empty set if a cycle of odd length exists in *AF*), and that it is not generally true that an *AF* has a unique preferred extension. In fact any *AF* that contains a cycle of even length may have multiple preferred extensions (see [10] for a proof). In the special case where there is a unique preferred extension we say the dispute is *resoluble*, since there is only one set of arguments capable of rational acceptance. Where there are multiple preferred extensions, we can view a *credulous* reasoner as one who accepts an argument if it is in *at least one* preferred extension, and a *sceptical* reasoner as one who accepts an argument only if it is in *all* preferred extensions.

If, however, we wish to allow that arguments may have different strengths, we open up the possibility that an attack can fail, since the attacked argument may be stronger than its attacker. Thus, if an argument attacks an argument whose value is preferred it can be accepted, and yet not defeat the argument it attacks. To represent this possibility of unsuccessful attacks we must extend the standard argumentation framework so as to include the notion of value.

To record the values associated with arguments we need to add to the standard argumentation framework a set of values, and a function to map arguments on to these values.

DEFINITION 2.2
A *value-based argumentation framework* (*VAF*) is a 5-tuple:
VAF = $<$*AR, attacks, V, val, P*$>$

where *AR*, and attacks are as for a standard argumentation framework, *V* is a non-empty set of values, *val* is a function which maps from elements of *AR* to elements of *V* and *P* is the set of possible audiences. The notion of value was introduced in Section 1 with the example of raising taxation being justified by the value of promoting equality, and lowering taxes by the value of stimulating enterprise. In general, an argument giving a case for performing an action will be justified relative to the end that the action is intended to secure, the purpose of performing the action, the value advanced by the action. In Section 4 we will propose an argument scheme which requires an explicit link between the action advocated and the value promoted for the argument to be well formed. We say that an argument *a* relates to value *v* if accepting a promotes or defends *v*: the value in question is given by *val(a)*. For every $a \in AR$, $val(a) \in V$.

The set *P* of audiences is introduced because, following Perelman, we want to be able to make use of the notion of an audience. We see audiences as individuated by their preferences between values, since if there is agreement on the ranking of values, there will be agreement on which attacks succeed.[1] We therefore have potentially as many audiences as there are orderings on *V*, and we can see the elements of *P* as being names for the possible orderings on *V*. Any given set of arguments will be assessed by an audience in accordance with its preferred values. We therefore next define an audience specific value-based argumentation framework, *AVAF*:

DEFINITION 2.3
An *audience specific value-based argumentation framework* (*AVAF*) is a 5-tuple:
$\text{VAF}_a = <AR, attacks, V, val, Valpref_a>$

where *AR, attacks, V* and *val* are as for a *VAF*, *a* is an audience, $a \in P$, and $Valpref_a$ is a preference relation (transitive, irreflexive and asymmetric) $Valpref_a \subseteq V \times V$, reflecting the value preferences of audience *a*. The *AVAF* relates to the *VAF* in that *AR, attacks, V* and *val* are identical, and *Valpref* is the set of preferences derivable from the ordering $a \in P$ in the *VAF*.

Our purpose in extending the *AF* is to allow us to distinguish between one argument attacking another, and that attack succeeding, so that the attacked argument is defeated. We therefore define the notion of *defeat for an audience*:

[1] Particular individuals are not permanently assigned to some audience. An individual must, for the purposes of a particular dispute be part of some audience, because there is an audience for every possible value order, and a consistent value order is a requirement of rationality. But individuals may change audiences from dispute to dispute, or even during a dispute. Individuals may even enter the dispute undecided as to their value order, and choose which to which audience they belong as the dispute proceeds, as in [11] and [21].

DEFINITION 2.4
An argument A $\in$ AF defeats$_a$ an argument B $\in$ AF for audience a if and only if both attacks(A,B) and not (val(B),val(A)) $\in$ Valpref$_a$.

Note that an attack succeeds if both arguments relate to the same value, or if no preference between the values has been defined. If *V* contains a single value, or no preferences are expressed, the *AVAF* becomes a standard *AF*. In practice we expect the number of values to be small relative to the number of arguments. Many practical disputes can in fact be naturally modelled using only two values. Note that defeat is only applicable to an *AVAF*: defeat is always *relative to a particular audience*. We write $defeats_a(A,B)$ to represent that *A* defeats *B* for audience *a*, that is *A* defeats *B* in VAF_a.

We next define the other notions associated with an AF for a VAF,

DEFINITION 2.5
An argument $A \in AR$ is *acceptable-to-audience-a* ($acceptable_a$) with respect to set of arguments *S*, ($acceptable_a(A,S)$) if:

$$(\forall x)((x \in AR \;\&\; defeats_a(x,A)) \rightarrow (\exists y)((y \in S) \;\&\; defeats_a(y,x))).$$

DEFINITION 2.6
A set *S* of arguments is *conflict-free-for-audience-a* if:

$$(\forall x)\,(\forall y)((x \in S \;\&\; y \in S) \rightarrow (\neg\, attacks(x,y) \vee valpref(val(y),val(x)) \in Valpref_a))).$$

DEFINITION 2.7
A *conflict-free-for-audience-a* set of arguments *S* is is *admissible-for-an-audience-a* if:

$$(\forall x)(x \in S \rightarrow acceptable_a(x,S)).$$

DEFINITION 2.8
A set of arguments *S* in a value-based argumentation framework *VAF* is a *preferred extension for-audience-a* ($preferred_a$) if it is a maximal (with respect to set inclusion) *admissible-for-audience-a* subset of *AR*.

We can illustrate the effects of these definitions by considering the argumentation framework containing four arguments, A, B, C and D, and two values, blue and red, shown in Figure 1.

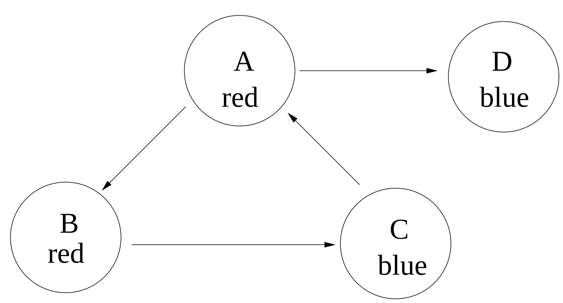

FIGURE 1. Value based argumentation framework.

Ignoring values and using the standard definitions of Dung, the preferred extension is the empty set. All arguments are attacked and do not defend themselves, and so none are acceptable by themselves.

Only two pairs of arguments, {B,D} and {C,D} are conflict free. B, however, does not attack A and so cannot defend D, and C cannot defend D from A either, as it is attacked by B, which is only attacked by A, which would introduce conflict if included.

Now suppose we see Figure 1 as a VAF with two values, either of which can be preferred, giving two audiences, Blue, which prefers blue to red, and Red, which prefers red to blue. Now, by Definition 2.5, D is acceptable to Blue by itself, since the attack of A on D fails through the value preference of this audience. Similarly C is acceptable to Blue because the attack by B fails given this preference. Moreover, the set {B,C,D} is conflict free for Blue by Definition 2.6, since the attack of B on C can be disregarded as blue is preferred to red. This set {B,C,D} is admissible for Blue, since the attacks of A on B and D are defended by C, which attacks A successfully for this audience. Finally {B,C,D} is a preferred extension for Blue, since no further arguments can be added to it.

Turning to the other audience, Red, A is acceptable to Red since the value preference of red over blue causes the attack of C to fail. A is in conflict for Red with both B and D, since its value is preferred, leaving {A, C} as a conflict free set for Red. This is also admissible, since the successful attack of B on C is defended by the successful attack of A on B. {B, D} is also conflict free for Red, but is not admissible for Red since there is no defence to the attacks of A on both these arguments. In a VAF therefore we have two preferred non-empty extensions: the preferred extension for Blue, {B,C,D} and the preferred extension for Red, {A,C}.

For a given choice of value preferences $Valpref_a$ we are able to construct an *AF* equivalent to the *AVAF*, by removing from *attacks* those attacks which fail because faced with a superior value.

Thus for any *AVAF*, $vaf_a = <AR, attacks, V, val, Valpref_a>$ there is a corresponding *AF*, $af_a = <AR, defeats>$, such that an element of *attacks*, *attacks(x,y)*, is an element of *defeats* if and only if $defeats_a(x,y)$. The preferred extension of af_a will contain the same arguments as vaf_a, the preferred extension for audience *a* of the *VAF*. Note in particular that if vaf_a does not contain any cycles in which all arguments pertain to the same value, af_a will contain no cycles. This is because the cycle will be broken at the point at which the attack is from an inferior value to a superior one for the audience concerned. Because multiple preferred extensions can only arise from even cycles, and empty preferred extensions only from odd cycles, both af_a and vaf_a will have a unique, non-empty, preferred extension for such cases. The *AFs* corresponding to the two *AVAFs* for the *VAF* in Figure 1 are shown in Figures 2 and 3.

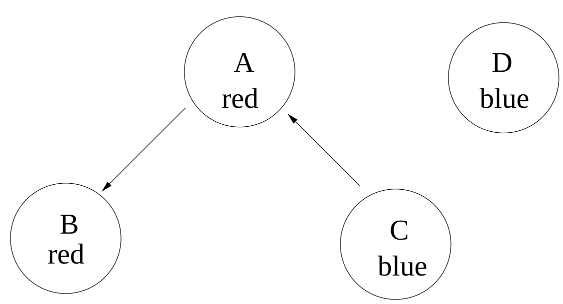

FIGURE 2. AF corresponding to AVAF$_{Blue}$.

Suppose we have a VAF with *k* values in *V*. *P* will contain factorial *k* distinct audiences, each of which will have a corresponding preferred extension $preferred_i$. Now for a given argument, A, in the VAF there are three possibilities:

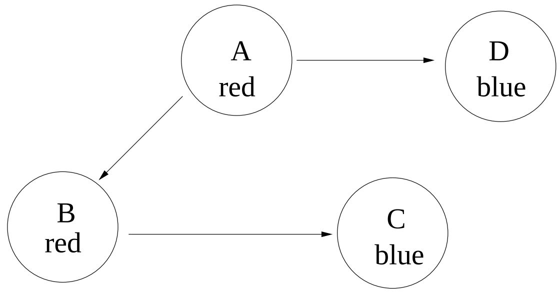

FIGURE 3. AF corresponding to AVAF$_{Red}$.

- A is in every *preferred*$_i$. In this case A is accepted by every audience, and so cannot be rejected by adopting a particular ranking of values. We therefore say that A is *objectively acceptable*.
- A is in at least one *preferred*$_i$, but not every *preferred*$_i$. In this case A will be acceptable to some audiences but not others, and so can either be accepted or rejected by choosing an appropriate value order. In this case we say that A is *subjectively acceptable*.
- A is in no *preferred*$_i$. In this case A will be rejected by every audience. We say that such an A is *indefensible*.

Although cycles do not give rise to multiple preferred extensions in VAFs, they remain interesting structures. Consider a cycle of length three in which the arguments relate to two values: suppose one argument is *blue* and the other two are *red*. In that cycle the blue argument will be objectively acceptable: if blue is preferred to red the attack on it will fail, and if red is preferred to blue the argument attacked by the blue argument will resist the attack and so defeat the attacker of the blue argument. The red arguments will both be subjectively acceptable: which is accepted will depend on the value preferences. In general, odd cycles with multiple values will always produce objectively acceptable arguments. Similarly, even cycles will, in certain configurations yield objectively acceptable arguments. For a full discussion and proofs see [10]. The significance of cycles in AFs and VAFs will become apparent in the next section when we discuss the representation of a particular body of case law.

3 Modelling case law as argumentation frameworks

In this section we will consider how a body of case law might be modelled in these frameworks. Modelling in terms of an AF was described in Bench-Capon [9], and values necessary to turn this AF into a VAF in Bench-Capon [11]. The body of case law used is the wild animal cases introduced to AI and Law by Berman and Hafner [15], and much discussed in recent years (e.g. Bench-Capon and Sartor [13], Bench-Capon and Rissland [12] and several papers in *Artificial Intelligence and Law* Volume 10 1-3). Here, as well as the three central cases described in Berman and Hafner [15], we will consider some additional cases. The facts of the chosen cases are:

Keeble v Hickergill (1707). This was an English case in which Keeble owned a duck pond, to which he lured ducks, which he shot and sold for consumption. Hickergill, out of malice, scared the ducks away by firing guns. The court found for Keeble.

Pierson v Post (1805). In this New York case, Post was hunting a fox with hounds. Pierson intercepted, killed and carried off the fox. The court found for Pierson.

Young v Hitchens (1844). In this English case, Young was a commercial fisherman who spread a net of 140 fathoms in open water. When the net was almost closed, Hitchens went through the gap, spread his net and caught the trapped fish. The case was decided for Hitchens.

Ghen v Rich (1881). In this Massachusetts case, Ghen was a whale hunter who harpooned a whale which subsequently was not reeled in, but was washed ashore. It was found by a man called Ellis, who sold it to Rich. According to local custom, Ellis should have reported his find, whereupon Ghen would have identified his lance and paid Ellis a fee. The court found for Ghen.

Conti v ASPCA (1974). In this New York case, Chester, a parrot owned by the ASPCA, escaped and was recaptured by Conti. The ASPCA found this out and reclaimed Chester from Conti. The court found that they were within their rights to do so.

New Mexico vs Morton (1975) and *Kleepe vs New Mexico (1976)*. These cases concerned the ownership of unbranded burros normally present on public lands, which had temporarily strayed off them. Both were won by the state.

The details of the modelling process need not concern us here: the interested reader is referred to Bench-Capon [9]. We may summarize the resulting model by saying that the cases were found to contain the twenty-four arguments given in Table 1, which could be related within an AF as shown in Figure 4.

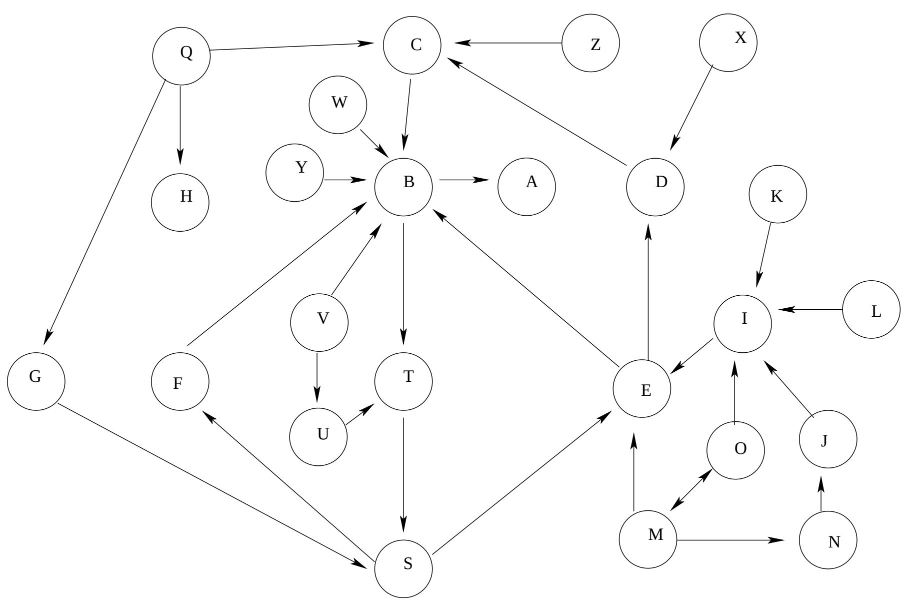

FIGURE 4. Argumentation framework for the animals cases.

TABLE 1. Arguments in the wild animal cases
CL = Clear law, UA = Useful Activity, PR = Protect property rights, EA = Economic Activity, CR = The Court Should Not Make Law

ID	Argument	Attacks	Values
A	Pursuer had right to animal		claim
B	Pursuer not in possession	A, T	CL
C	Owns the land so possesses animals	B	PR
D	Animals not confined by owner	C	
E	Effort promising success made to secure animal made by pursuer	B, D	CL
F	Pursuer has right to pursue livelihood	B	EA
G	Interferer was trespassing	S	PR
H	Pursuer was trespassing	F	PR
I	Pursuit not enough (JUSTINIAN)	E	CL
J	Animal was taken (JUSTINIAN)	I	CL
K	Animal was mortally wounded (Puffendorf)	I	CL
L	Bodily seizure is not necessary (Barbeyrac), interpreted as animal was brought within certain control (TOMPKINS)	I	UA
M	Mere pursuit is not enough(TOMPKINS)	E, O, N	CL
N	Justinian is too old an authority (LIVINGSTON)	J	
O	Bodily seizure is not necessary (Barbeyrac), interpreted as reasonable prospect of capture is enough (LIVINGSTON)	I, M	UA
Q	The land was open	G, H, C	PR
S	Defendant in competition with the plaintiff	E, F	EA
T	Competition was unfair	S	EA
U	Not for courts to regulate competition	T	CR
V	The iron holds the whale is an established convention of whaling	B, U	CR
W	Owners of domesticated animals have a right to regain possession	B	PR
X	Unbranded animals living on land belong to owner of land	D	PR
Y	Branding establishes title	B	
Z	Physical presence (straying) insufficient to confer title on owner	C	CL

This model could be used to consider a given case by first removing the nodes representing arguments not applicable in the case under consideration, and any attacks they make. For example, in *Pierson*, where the land was open, neither G nor H represents an applicable argument. Next the Preferred Extensions of the resulting pruned framework can be determined. If there is a unique preferred extension, then the case is resolved for the plaintiff if it contains A and for the defendant if it does not. The AF in Figure 4, however, contains two even length cycles:

- the two cycle M-O, which arises in Pierson
- the four cycle B-T-S-E, which arises in Young

In these cases there will be no unique preferred extension, but rather two preferred extensions, one with and one without A. Thus representation as an AF can draw attention to the key point of contention: in the case of *Pierson*, whether or not the prospect of capture is enough to attribute possession to the pursuer, and in the case of Young, whether we accept the pair of arguments that the defendant was in competition with the plaintiff who did not possess the fish, or the pair that the plaintiff was in furtherance of his livelihood and the competition was unfair.

Although use of a Dung-style Argumentation Framework can isolate points of contention in particular cases, it does not provide any basis for resolving such dilemmas. One basis was proposed by Berman and Hafner in [15], which says that such dilemmas should be resolved through consideration of the purposes that the law is intended to effect, the values it is intended to promote. It is this suggestion that the subsequent work cited in the first paragraph of this section has explored. Value Based Argumentation Frameworks [10] were motivated by a desire to provide a formal basis for this notion, by extending Dung's framework to handle values.

VAFs have been applied to the animals problems in [11]. The values associated with the arguments are shown in the third column of Table 1. Two things emerge from construing the representation as a VAF. First we can see that some lines of argument are not worth pursuing. If B is attacked by C, there is no point in attacking C with Z, since if Clear Law is upheld over Property Rights, the attack on C is not needed, and if Property Rights are preferred to Clear Law it cannot succeed. Thus the use of values may help to prune certain lines of argument.

Second it helps with the cycles. In the case of the two cycle, we can see that whether we accept M or O turns on whether we believe that promoting the useful activity is worth the resulting unclarity in the law. The case of the four cycle is still more interesting. The four cycle B–T–S–E involves two arguments of one value followed by two arguments of another, and we know from [10] that this means that the first two arguments in each value, T and E in this case, are *objectively* acceptable. This means that if we accept that the competition was unfair, we must hold that Young's efforts entitled him to the fish, and find for the plaintiff. Unlike the two cycle this is not a matter of choice, it holds however we order our values.

If the defendant is to win the case, he must find a way to break the cycle. This was done in *Young* by arguing that it was for the legislature, not the courts, to decide what constituted unfair competition. Note that it could not be argued that the competition was not unfair: all intuitions seem to suggest that it was indeed unfair, although breaking no written law. Now by giving priority to this principle over Economic Activity, T can be defeated and the cycle broken. Note that V, as used in Ghen, is still able to defeat U: in upholding a convention universally respected within an industry the court is not making a law at its own discretion.

In this section we have discussed the representation of case law as Argumentation Frameworks and Value Based Argumentation Frameworks. This provides a useful way to identify points of disagreement which can arise and can provide some help in resolving them. Both, however, required a purely manual modelling of the domain to produce the frameworks. In the next section we will discuss how we can generate VAFs in the context of a multi-agent system.

4 Argument schemes for persuasion about action

The study of AI has produced much research on the topic of reasoning about beliefs. However, reasoning about action — practical reasoning — has received comparatively little attention in Computer Science. One way to approach the topic of practical reasoning is to treat it as a species of presumptive argument through the use of argument schemes and associated critical questions. This approach has been advocated by Walton in [27]. His account uses an argument scheme to put forward a presumptive justification for performing an action based on the action bringing about some desired goal. The presumption behind the argument can then be subjected to a series of critical questions. Whether this presumption stands or falls then depends on satisfactory answers being given to any of these critical questions posed in the particular situation. Subjecting the argument to such challenges is how we hope to identify any appropriate alternatives that require consideration. This will then enable us to determine the best choice of action to be taken, in the particular context.

In [5] we extend Walton's *Sufficient Condition Scheme for Practical Reasoning* to make it more precise and this has led us to propose the following argument scheme for the justification of an action:

AS1 In the circumstances R
we should perform action A
to achieve new circumstances S
which will realize some goal G
which will promote some value V.

Representing the justification in this manner enables us to separate out parts of the argument which conflict and this allows us to identify the precise part of the position that an opponent may disagree with. More specifically, this scheme makes the notion of a goal more precise by showing that an action may be justified in terms of: its direct consequences (the effects of the action); a state of affairs following from the direct consequences which the action was intended to realize (the goal of the action); the underlying social value promoted by performing the action (the purpose of the action).

In order to challenge the presumptions present in this argument scheme we go on to specify sixteen critical questions associated with it (compared with the four identified by Walton). They are as follows:

- CQ1: Are the believed circumstances true?
- CQ2: Assuming this, does the action have the stated consequences?
- CQ3: Assuming all of these, will the action bring about the desired goal?
- CQ4: Does the goal realize the value intended?
- CQ5: Are there alternative ways of realizing the same consequences?
- CQ6: Are there alternative ways of realizing the same goal?
- CQ7: Are there alternative ways of promoting the same values?
- CQ8: Does doing the action have a side effect which demotes the value?
- CQ9: Does doing the action have a side effect which demotes some other value?
- CQ10: Would doing the action promote some other value?
- CQ11: Does doing the action preclude some other action which would promote some other value?
- CQ12: Is the action possible?
- CQ13: Are the circumstances as described possible?
- CQ14: Are the consequences as described possible?
- CQ15: Can the desired features be realised?

- CQ16: Is the value indeed a legitimate value?

Each of the above critical questions represents a source of disagreement and in posing a critical question an opponent is making what we refer to as an 'attack' on the element of the position in question. This attack may be stated with a varying degree of force and thus this leads to a number of variants that can be associated with particular critical questions. For example, examining CQ1: an attacker may simply express disagreement with some aspect of a position, as when an attacker denies that R is the current state of the world. Beyond this minimalist attack, an attacker may also state an alternative position to that proposed, for example, expressing not only that R is not the current state of the world, but instead that T is the current state. A full list and description of the attacks and their variants are given in [7, 23].

Additionally, each critical question falls into one of three distinct categories which relate to the nature of the attack on the presumption: issues relating to the beliefs as to what is the case; issues relating to desires as to what should be the case; and issues relating to representation concerning the language being used and the logic being deployed in the argument. In [7] there is a detailed discussion of each of these categories into which the critical questions fall. Resolution of conflicting points is dependent upon the category into which the attack falls. Disputes over facts should be straightforwardly resolved if some process of empirical investigation is agreed upon between the participants, or if the question can be put to some decisive arbiter such as a jury. Issues of representation should also be capable of resolution by agreeing on language and context before the dialogue starts, and by aligning the ontologies of the participants to ensure a shared understanding of the concepts in the given topic of conversation. Both disagreements about representation and disagreements about facts should be resolved before disagreements about choice can be addressed. Resolution of disputes about what is best typically depends on the context in which the dialogue is taking place. It may be the case that one party is an authority on the matter in question and so this will facilitate resolution. For example, in government issues it is usual for government advisors to find out the facts of the situation, and then ministers make the choices between actions on the basis of these facts. Naturally, resolution will also occur if one party allows himself to be persuaded that his preference ordering is wrong or to accept the ordering of his opponent's preferences. Disagreements turning on values require extra attention as it is often difficult to get an opponent to change his values. In order to facilitate and resolve such issues we can make use of VAFs, as discussed further below.

In [7, 4] we make the above account of practical reasoning computational. Conditions under which BDI agents can put forward a position of the form stated above are specified and preconditions which need to be satisfied in order for an opposing agent to make an attack on this position are given. This requires an addition to the standard BDI model by including a mechanism to allow agents to make use of value functions, as found in our argument scheme for practical reasoning. This extension is made by associating each desire with a value, which gives the reason why it is desirable.

This computational account enables BDI agents to reason and argue about which action should be taken in a particular context and we have applied our account to a number of different domains, including medicine and politics as well as law. With regard to the legal application, in [6] we show how the reasoning in the well known property law case of *Pierson vs Post*, 3 Cai R 175 2 Am Dec 264 (Supreme Court of New York, 1805) can be reconstructed in terms of our account. In this example our account allows us to model the various participants in the debate as different agents. We see the disagreements as grounded in divergent beliefs, goals and values, and therefore use different agents to represent the different views that can be brought to bear on the problem. We use four different agents named A, B, L and T, to represent the different beliefs, desires and values pertinent to the problem and generate the arguments that these agents can form on the basis of this knowledge. Broadly A represents the view that fox hunting should be protected as a sport, B that fox hunting should be abolished because it is cruel, L the minority opinion of Livingston, that fox hunting should be protected because of its social utility in destroying foxes, and T the majority opinion of Tompkins that ownership of a wild animal can be established only through possession. Based upon the beliefs and desires to which each individual agent subscribes, they can each provide one or more instantiation of AS1. This enables us to show the relations between these arguments as a set of VAFs and use the procedures for resolving conflicts in VAFs, as described earlier in Section 2, to calculate the dialectical status of the arguments with respect to the different audiences represented by the different agents. As mentioned previously, the notion of audience preferences is crucial as persuasion to accept an argument relies on the value preferences of the audience rather than the speaker. The speaker can hope that the audience will recognize some desirable consequences and acceptable values for themselves and this is especially important in the legal domain.

In reconstructing the argumentation of *Pierson* we use three separate argumentation frameworks to show the views of the agents at three different levels: the level of facts about the world, at which desires are derived; the level at which the legal system connects with the world to achieve these desires; and at the level of pure legal concepts. These levels are familiar from other work in AI and Law, and are explicit in the functional ontology of Valente [26], and some discussions of expert systems within the logic programming paradigm, such as [8]. Conclusions at lower levels will be used as premises at higher levels. We present each of the argumentation frameworks showing the attacks made on the instantiations of AS1. In the figures below, nodes represent arguments. They are labelled with an identifier, the associated value (which can be determined from the argument scheme represented by the node), if any, and on the right-hand side, the agent introducing the argument. Arcs are labelled with the number of the attack they represent from our theory of persuasion over action [23]. We then summarize what can be deduced from the framework in order to proceed to the next level in the argument. Figure 5 shows the argumentation framework for level 1.

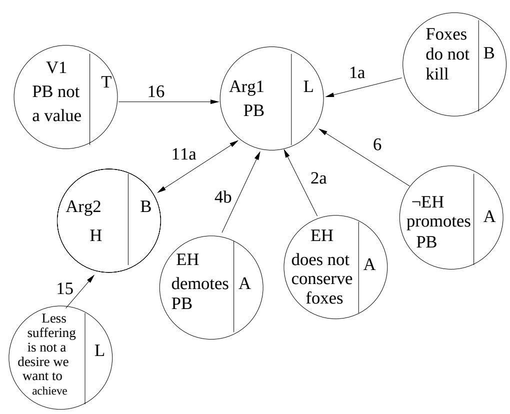

FIGURE 5. Level 1: Arguments about the world. (Goals: EH = Encourage Hunting. Values: PB = Public Benefit, H = Humaneness.)

The following instantiations of AS1 are used in this framework. We omit S, the circumstances resulting from the performance of the action since G represents the relevant subset of these circumstances. S is of importance only if we need to distinguish what results from an action from the desires that it satisfies. Attacks on these schemes are shown in the Figures 5–7 as labels on the arcs:

Arg1

R1: Where foxes kill livestock, encouraging hunting leads to fewer foxes and fewer foxes means farmers are protected
A1: encourage fox hunting
G1: as fewer foxes and farmers protected
V1: promotes public benefit.

Arg2

R2: Where fox hunting is cruel
A2: discourage fox hunting
G2: as reduced animal suffering
V2: promotes humaneness.

From the VAF in Figure 5 we can now deduce the arguments that each agent finds acceptable. Agent L can, by making suitable choices about preferences, deduce that hunting should be encouraged and agent B that hunting should be discouraged, using his own preferences. A and T, by accepting L's argument against Arg2, need subscribe to neither argument, and so derive no additional desires from this level of the debate. We can now move on to the next level, giving the argumentation framework shown below in Figure 6.

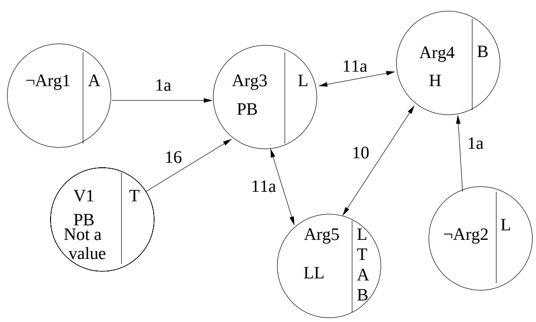

FIGURE 6. Level 2: Linking to legal concepts. (Values: PB = Public Benefit, H = Humaneness, LL = Less Litigation.)

This argumentation framework is constructed from the following instantiations of AS1:

Arg3

R3: Where there is pursuit and fox hunting is to be encouraged
A3: find ownership established
G3: as hunting encouraged
V3: promotes public benefit.

Arg4

R4: Where there is pursuit and fox hunting is to be discouraged
A4: find ownership not established
G4: as hunting discouraged
V4: promotes humaneness.

Arg5

R5: Where there is no possession
A5: find ownership not established
G5: as finding no ownership where no possession
V5: promotes less litigation.

In this second framework we can see that all agents except L agree that Arg3 is defeated, although for different reasons, and so accept Arg5. L accepts Arg5, but believes that its force is insufficient to defeat Arg3 since he prefers Public Benefit to Less Litigation. We can now move on to the top level arguments, giving the argumentation framework shown below in Figure 7:

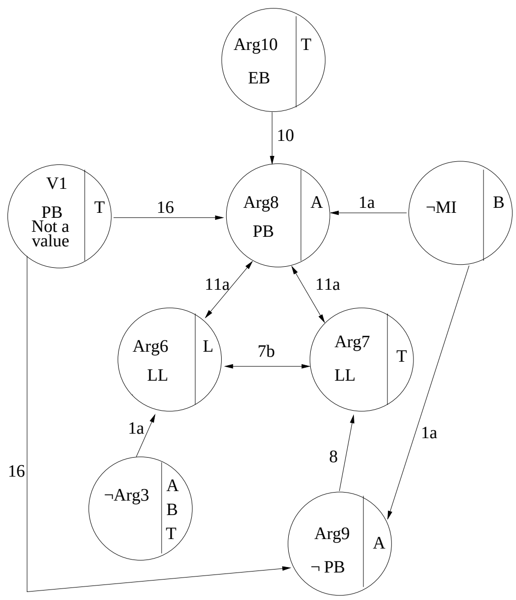

FIGURE 7. Level 3: Arguments in terms of legal concepts. (Values: PB = Public Benefit, H = Humaneness, LL = Less Litigation, MI = Malicious Intent, EB = Economic Benefit).

Arg6

R6: Where there is ownership
A6: find for plaintiff
G6: as finding for plaintiff with ownership
V6: promotes less litigation.

Arg7

R7: Where there is no ownership
A7: find for defendant
G7: as finding for defendant where there is no ownership
V7: promotes less litigation.

Arg8

R8: Where there is malicious interference by defendant
A8: find for plaintiff
G8: as finding for plaintiff where there is malicious interference
V8: discourages immoral behaviour.

Arg9

R9: Where there is malicious interference by defendant
A9: do not find for defendant
G9: as finding for defendant where there is malicious interference
V9: encourages immoral behaviour.

Arg10

R10: Given the facts of Keeble
A10: do not find for defendant
G10: as finding for defendant where there is malicious interference and productive activity
V10: demotes economic benefit.

This now completes the final framework and so we can deduce whether the plaintiff has remedy or not. L, who accepts Arg3, and gives prime importance to public benefit, will use Arg6 to determine his decision. A, who also gives primacy to public benefit, but rejects the facts on which Arg6 is ultimately based will use Arg8. B rejects both Arg6 and Arg8 on factual grounds, and so accepts Arg7, and finally T accepts Arg7 as it is the only argument grounded on a value of which the law should take note. Arg10 illustrates the use of precedents as described in [23]. This completes the analysis of the different views of the agents represented in the case. In [6] we go on to discuss the decision made in the actual case of *Pierson vs. Post* and we relate the opinions given in the case to the various components of the VAFs produced in our computational example as given above. If we refer back to the representation of *Pierson* in Section 3 and depicted in Figure 4, we can see that the two cycle M–O appears as the two cycle Arg3–Arg5 in Figure 6, reflecting the value choice between public benefit and clear law.

We believe that this example shows how we can generate VAFs to reconstruct the reasoning in legal cases. It shows the feasibility of generating and construing legal argument in terms of our argument scheme and critical questions, and how VAFs support decisions regarding these arguments. We have shown that this model can accurately represent the reasoning deployed in such cases and how we can realize such debates as a multi agent system, with the different agents representing different audiences with divergent beliefs, desires and values.

5 Construction of theories from case law

In the previous section we showed how VAFs can be generated by agents who have information linking situations and actions and values. In this section we discuss how this kind of knowledge may be produced in a case law domain. Our starting point is the account of Bench-Capon and Sartor [14] which views reasoning with legal cases as the construction of a theory, based on precedent cases, which links facts and values and reconciles conflicting rules with value preferences. We describe a set of tools which support the construction of such theories: CATE, which is a tool for building such theories manually; and AGATHA which automatically constructs a space of theories by using a sequences of arguments moves. AGATHA can be used to produce the entire theory space if the number of precedents is sufficiently restricted to allow this, or, coupled with a means of evaluating theories (ETHEL), uses heuristics to guide its search.

5.1 Construction of theories from cases

In [14] Bench-Capon and Sartor take as a starting point a background comprising a set of cases (‘the case background’) each described using some set of descriptors (‘the factor background’). In their initial presentation the descriptors correspond to factors as used in CATO [1], although [14] also presents extensions allowing cases to be described in terms of dimensions as used in HYPO [2]. The factor background construes each descriptor as a reason for deciding the case for one of the parties, either the plaintiff or the defendant, and associates each descriptor with this party and a value. Values represent the social value enhanced or promoted by deciding the case for the favoured party given the presence of this reason. Thus, for example, in Trade Secrets Law, that the plaintiff took security measures is a reason for deciding for the plaintiff, and promotes the value that fairness requires that people should take reasonable steps to protect their own interests.

The use of values is important in [14]. The idea, originally proposed by Berman and Hafner in [15], is that cases reveal preferences between conflicting reasons, and that these preferences are explained in the theory through the priority given to the values promoted by these reasons. Values, in this way, allow us to deduce preferences which have not yet been explicitly tested in any precedent case. Consistency is required of preferences between rules and values, since this is an essential condition of any coherent theory.

Against this background, the aim of theory construction when presented with some new cases is to produce a theory which provides a sufficient explanation of the previous cases, and which gives an answer in the new case. Competing theories may be constructed, and these need to be critically compared to determine which should be applied to the new case.

A theory itself comprises a five-tuple:

- a set of cases selected from the case background (C);
- a set of factors selected from the factor background (F);
- a set of rules, comprising simple rules relating factors to the side favoured, and composite rules with several reasons in the antecedent (R);
- a set of preferences between these rules (RP);
- a set of preferences between the various values promoted by the factors (VP).

Each of these five sets is initially empty. The theory is constructed by adding to them using one of a number of theory constructors defined in [14]. Here we give informal explanations of the theory constructors: formal definitions are given in [14].

- *Include Case*: This adds a case from the case background to C.
- *Include Factor*: This adds a factor from the factor background to F. Note that the cases in the theory are described only using the factors which have been included. Thus adding a factor to F may modify the descriptions of the cases in C. Additionally the factor adds to R a rule expressing that the factor is a (defeasible) reason to decide for the party it favours.
- *Factors Merging*: Given a rule in R, the antecedent may be strengthened to give a new rule. Antecedents may be strengthened only by the addition of another factor favouring the same party to the dispute.

- *Role Broadening*: Given a rule in R, the antecedent may be weakened to give a new rule by omitting one of the factors from the antecedent.
- *Rule Preference from Cases*: Given a case in C to which two rules in R each favouring a different party are applicable, we may infer a preference for the party which won the case and add this to RP. Moreover, from this rule preference we may infer that the set of values promoted by following the preferred rule are preferred to those promoted by following the other rule, and add this to VP.
- *Rule Preference from Value Preference*: Given a preference in VP and two rules corresponding to the related sets of values, we may deduce that the rule relating to the preferred value is preferred to the other rule, and include this in RP.
- *Arbitrary Rule Preference*: Using this constructor a preference is added to RP, even though no case can be found to justify it.
- *Arbitrary Value Preference*: Using this constructor a preference is added to VP, even though no rule preference can be found to justify it.

5.2 CATE

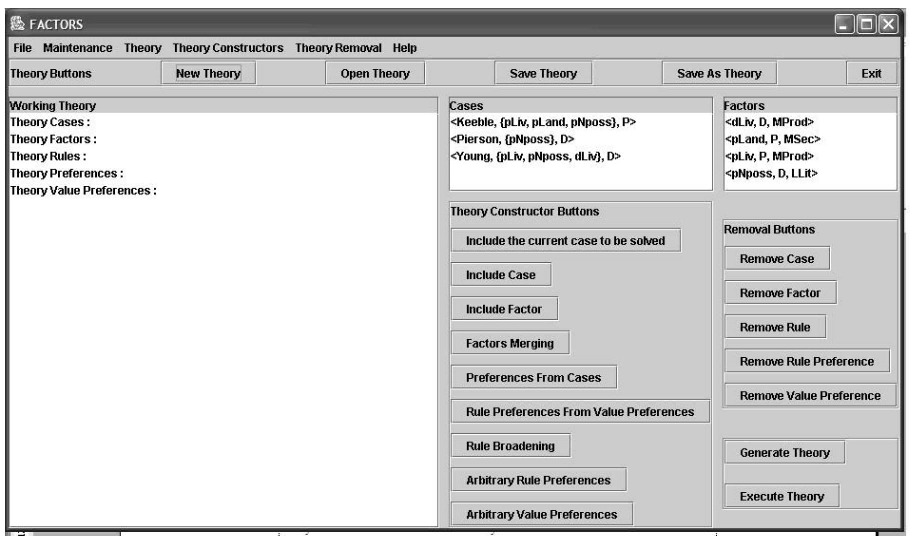

FIGURE 8. The CATE theory construction tool.

CATE, shown in Figure 8, is designed to embody the set of theory constructors described in [14] and has been implemented in Java. There are panels to show case and factor backgrounds, and the theory under construction. The various theory constructors can be used by clicking on the appropriate button on the screen. For example, to include a case into the theory, the *Include Case* button is selected and the user is prompted to choose which case they want. CATE also provides some checks on the legality of use of the constructors: when a user specifies preferences over rules or values, CATE checks that the resulting theory is consistent and, if adding the preference would make the theory inconsistent, then a warning is issued and the preference is not added. If the user of CATE still wishes to include the preference then the existing preference causing the conflict must first be removed. CATE also tracks where the rule preferences came from, so that we can distinguish those derived from cases from those derived from a value preference, by labelling each rule preference depending on which theory constructor was used.

The theory can be executed by clicking on a button and CATE will generate the Prolog code representing the theory which can then be executed to give the outcome of applying the theory to each case included in it. These case outcomes can then be used to evaluate how the theory performs with respect to the actual decisions for the cases, so as to verify that the theory does indeed explain the selected cases.

When the theory is executed it is analysed to produce executable Prolog code by the following steps. First the value preferences are converted into the set of corresponding rule preferences. This is done by converting each value in the preference into its associated factors. The tool checks for any inconsistencies and only if the theory is consistent will the code generation continue.

Now the rule preferences are used to order the rules in the theory. The rules are held in a list and this list is compared to the rule preferences. The list is thus guaranteed to be ordered in conformance with the rule preferences in the theory, but, since these preferences are not complete, this does not determine a unique ordering. Conflicts are resolved using the alphabetical order of the rules contained in the theory.

Finally the Prolog code is generated and output to a file that can be executed in a standard Prolog interpreter or by CATE itself. The outcomes for every case included in the theory are recorded in a text file which can be analysed or exported to a spreadsheet package and manipulated.

Several experiments were performed using CATE to test various ideas of how to construct and compare theories, the results of which are described fully in [18]. Broadly the experiments confirmed that it was possible to use values in the way suggested by Bench-Capon and Sartor, that all the available factors should be taken into account, and that giving values different weights improved the explanatory power of the theories.

5.3 AGATHA

AGATHA models adversarial dialogue between two agents with each agent taking turns to make a move to adapt a theory so that it produces an outcome for their side. As described below, AGATHA has five moves that it can use according to certain preconditions and it applies all possible moves at each point until no more moves can be made to generate the complete theory space.

AGATHA checks which moves can be made by checking the preconditions for each move against the theory at that point in the game tree and, if the preconditions match, it applies the move. Each move that can be applied produces a new theory. When alternative moves are available, new branches are added to the tree of theories being created. As each move is applied to the theory, the resulting theories are examined and only those which give the same outcome for the problem case as the party making the move are retained. If the move made does not give the desired outcome, the theory is discarded because, even though the move could be applied, it does not help the party making the move, and so does not represent a sensible move.

The five moves available in AGATHA are described fully in [19]: here we give a short description of these moves:

1. *Analogize Case*. This move cites a precedent case which has the outcome the party making the move desires, and uses the factors in common with the current case.

2. *Distinguish with Case*. This move distinguishes a case already cited in the debate and cites a new case which has a different outcome.
3. *Distinguish with Arbitrary Preference*. This move distinguishes the previously cited case in the same way as for the *Distinguish with Case* move, but instead of analogizing a new case, an arbitrary preference for the set of factors favouring the winning side in the precedent case over the set of factors favouring the losing side in that case is expressed. Assuming that the new case lacks a factor present in the original case, the preference from the precedent is no longer applicable to the new case.
4. *Distinguish Problem*. This move distinguishes the problem case instead of the previously cited case. That is, it expresses a preference for the factors in the current case favouring the side which lost the precedent case over the factors in the new case favouring the other side. Assuming that the new case contains a factor not present in the original case, the preference cannot then be applied to the precedent case.
5. *Counter with Case*. This move counters the previously cited case by finding a case which is as-on-point or more-on-point as the previous case but was decided for the other side.

The argument moves used in AGATHA use the same theory constructors used in CATE to create the underlying theory, and to generate executable code. Each move corresponds to a number of theory constructors which are applied to extend the current theory.

Again experiments have been carried out to show that AGATHA can construct useful theories. Among them, AGATHA has used the wild animals cases described in Section 3, and reproduced the sequence of theories produced by hand in [14]. Moreover, by using these moves, which correspond to the argument moves of Case Based Reasoning systems such as HYPO and CATO, AGATHA is following a cognitively plausible strategy, and the sequence of moves used to reach a theory is available to present the case to an opponent. The results are described fully in [19].

If there is a reasonable number of cases in the case background, the tree representing the theory space rapidly becomes very large. We therefore wish to use heuristic evaluation of the theories to guide our expansion of the tree. This requires a way of evaluating theories so as to choose which branches to develop. Section 5.4 describes how we evaluate the theories.

To provide this heuristic search AGATHA was extended to provide a cooperative search heuristic based on A* and is currently being extended to provide an adversarial search heuristic based on Minimax or $\alpha\beta$ pruning. A* is not an adversarial search, and so using it in AGATHA makes theory construction essentially a cooperative process, despite the alternation of plaintiff and defendant theories. No account is taken of how good a response a move permits, and so no notion of blocking the 'opponent's' best moves can be used.

5.4 ETHEL

ETHEL stands for Evaluation of THEories in Law and evaluates theories using criteria similar to those proposed in [14], including explanatory power, simplicity, freedom from arbitrary preferences and the ability to generalise to new cases.

Each theory is evaluated using the following five criteria to give an Evaluation Number which can be used to compare the theories.

1. *Simplicity*. ETHEL counts the number of rules in the theory, the number of arbitrary rule preferences and the number of rule preferences obtained from value preferences. A simpler theory is better than a more complex theory. The simplest theory would only contain one rule preference

and this should not be an arbitrary rule preference or a rule preference from value preference. The theory is penalised if it contains arbitrary rule or value preferences.

2. *Explanatory Power*. Each theory is executed with the complete set of background cases and the results analysed. The value for the Explanatory Power is given by the number of correctly decided cases plus half the abstention cases divided by the total number of cases and multiplied by 100 to give a percentage of the total number of cases.
3. *Completion Explanatory Power*. The previous criterion executes a program which uses only the factors specifically used in the construction of the theory. This gives a restricted set of factors to be considered when deciding the cases. For this criterion the unused background factors are loaded into the theory, and preferences between rules using these factors are inferred on the basis of the value preferences in the theory. This extended theory is executed on the complete set of background cases. This number is calculated in the same way as for the Explanatory Power.

   The above three numbers are summed to give a basic Evaluation Number which is based on how well the theory performs in explaining the background cases and its simplicity. We now adjust this number according to the position of the theory in the tree.
4. *Depth*. The basic Evaluation Number can be increased by adding an amount which represents how deep the theory is in the tree. This encourages AGATHA to explore the search space more deeply.
5. *Leaf Node*. The depth-extended Evaluation Number can be increased again by adding an amount which represents whether or not the theory is a leaf theory, to reflect the fact that this theory cannot be profitably modified by an opponent.

These Evaluation Numbers give a value with which to compare the theories based on how well they explain the background, their structure and their position in the development of the game tree. They can be used to evaluate the nodes in the theory tree, and so guide a heuristic search. Our heuristic search is based on A*, perhaps the best known such algorithm, described in most standard AI textbooks (e.g. [28]).

A search heuristic such as A* is necessary if we are to make full use of available background cases. The ability to use a more extensive background does improve the results for AGATHA over those obtained from generating the complete theory space on a limited background. Moreover, as reported in [19] our experiments have shown that AGATHA and ETHEL together produce better theories than the hand constructed theories reported in [17]. Indeed, the theories produced are comparable in explanatory power to the best performing reported technique, IBP [16], [3], achieving an accuracy of over 90%.

6 Concluding remarks

In the preceding sections we have shown how we can represent reasoning with legal cases in terms of argumentation frameworks, the improvements that result from associating values with arguments, how we can generate VAFs and how we can discover the knowledge required to generate VAFs from precedents. A key question remains, however, as to the nature of these values. In some contexts they look like legal principles (e.g. Clear Law), while in other contexts they represent behaviour which it is wished to encourage (e.g. the values used in CATE and AGATHA), and in yet other contexts general moral positions (e.g. humaneness). In [11] it was suggested that values could arise from a number of sources. First they can appear as the principles governing a particular juridical culture (for a full discussion of such principles and how they vary across jurisdictions, see [20]). Second they can arise from prevailing social consciousness: such values can be volatile and can motivate changes in law. The development of the treatment of women in law in the UK over the twentieth century is a good example of this. Third they can appear from ideology: the left–right debate of the twentieth century providing an example of this. As for preferences, the judge can be seen as an oracle that pronounces on what is to be preferred.

The discussion in Section 4 of this paper offers a different, but complementary, perspective, and some explanation of why values can cover such a variety of things. At the lowest level we find instrumental values — social values that determine which behaviours it is felt should be encouraged by the law. Such behaviours can then be used to ground arguments about the qualification of legal concepts at the second level. It is at this level that much of the reasoning found in traditional case based reasoning such as CATO goes on, and this explains why desired behaviours are the natural choice for values in CATE and AGATHA. At the highest level legal principles take over, using the legal qualifications established at the second level. In particular cases, as *Pierson* illustrates, we may need to consider arguments relating all three kinds of value, and the question of where values come from can only be properly answered in terms of the stage of the reasoning.

Practical reasoning has been found to have a number of distinctive characteristics which any account of practical reasoning should be able to accommodate. These include:

- Practical reasoning is defeasible, and requires the consideration of alternatives and cons as well as pros. For this reason such arguments must always be considered in the context of attacking and supporting arguments.
- Practical reasoning concerns choice, and must allow for the possibility of rational disagreement. For this reason we must handle different audiences accepting different sets of arguments.

Legal argumentation shares these features with practical reasoning, which is why it is illuminating to construe legal reasoning in terms of an account of practical reasoning, as we have done in this paper. The need for context motivates the use of argumentation frameworks, and the need to distinguish audiences motivates the use of values. Moreover we have demonstrated how we can reconstruct legal argumentation using this approach, and how we can derive the required knowledge from precedent cases. Much remains to be explored, but we hope that we have provided some foundations on which to build.

Acknowledgements

This paper is based on a talk given by the first author at the 6th de Morgan Workshop, Logic and Law, held at King's College, London, in December 2004.

Katie Atkinson and Alison Chorley are grateful for support from the EPSRC. Thanks also to Paul Dunne for many fruitful discussions concerning formal argumentation frameworks.

References

[1] V. Aleven. *Teaching Case Based Argumentation Through an Example and Models*. Phd thesis, University of Pittsburgh, Pittsburgh, PA, USA, 1997.

[2] K. D. Ashley. *Modeling Legal Argument*. MIT Press, Cambridge, MA, USA, 1990.

[3] K. D. Ashley and S. Brüninghaus. A predictive role for intermediate legal concepts. In *Proceedings of the Sixteenth Annual Conference on Legal Knowledge and Information Systems (Jurix 2003)*, D. Bourcier, ed., pp. 153–162, Amsterdam, The Netherlands, 2003. IOS Press.

[4] K. Atkinson, T. J. M. Bench-Capon, and P. McBurney. Attacks on a presumptive argument scheme in multi-agent systems: pre-conditions in terms of beliefs and desires. Technical Report ULCS-04-015, Department of Computer Science, University of Liverpool, UK, 2004.

[5] K. Atkinson, T. J. M. Bench-Capon, and P. McBurney. Justifying practical reasoning. In *Proceedings of the Fourth International Workshop on Computational Models of Natural Argument (CMNA 2004)*, F. Grasso, C. Reed, and G. Carenini, eds, pp. 87–90, Valencia, Spain, 2004.

[6] K. Atkinson, T. J. M. Bench-Capon, and P. McBurney. Arguing about cases as practical reasoning. In *Proceedings of the Tenth International Conference on Artificial Intelligence and Law (ICAIL 2005)*, pp. 35–44, New York, 2005. ACM Press.

[7] K. Atkinson, T. J. M. Bench-Capon, and P. McBurney. Computational representation of practical argument. *Knowledge, Rationality and Action*, 2005. *In press*.

[8] T. J. M. Bench-Capon. Knowledge based systems applied to law: A framework for discussion. In *Knowledge Based Systems and Legal Applications*, T. Bench-Capon, ed. pp. 329–342, Academic Press, London, 1991.

[9] T. J. M. Bench-Capon. Representation of case law as an argumentation framework. In *Proceedings of the Fifteenth Annual Conference on Legal Knowledge and Information Systems (Jurix 2002)*, T. J. M. Bench-Capon, A. Daskalopoulu, and R. Winkels, eds, pp. 103–112, Amsterdman, The Netherlands, 2002. IOS Press.

[10] T. J. M. Bench-Capon. Persuasion in practical argument using value based argumentation frameworks. *Journal of Logic and Computation*, **13**, 429–48, 2003.

[11] T. J. M. Bench-Capon. Try to see it my way: Modelling persuasion in legal discourse. *Artificial Intelligence and Law*, **11**, 271–87, 2003.

[12] T. J. M. Bench-Capon and E. L. Rissland. Back to the future: dimensions revisited. In *Proceedings of JURIX 2001*, B. Verheij, A. Lodder, R. Loui, and A. Muntjewerff, eds, pp. 41–52, Amsterdman, The Netherlands, 2001. IOS Press.

[13] T. J. M. Bench-Capon and G. Sartor. Theory based explanation of case law domains. In *Proceedings of the Eighth International Conference on AI and Law (ICAIL 2001)*, pp. 12–21, New York, 2001. ACM Press.

[14] T. J. M. Bench-Capon and G. Sartor. A model of legal reasoning with cases incorporating theories and values. *Artificial Intelligence*, **150**, 97–143, 2003.

[15] D. H. Berman and C. D. Hafner. Representing teleological structure in case-based legal reasoning: the missing link. In *Proceedings of the Fourth International Conference on AI and Law (ICAIL 1993)*, pp. 50–59, New York, NY, USA, 1993. ACM Press.

[16] S. Brüninghaus and K. D. Ashley. Predicting outcomes of case-based legal arguments. In *Proceedings of the Ninth International Conference on AI and Law (ICAIL 2003)*, pp. 233–42, New York, NY, USA, 2003. ACM Press.

[17] A. H. Chorley and T. Bench-Capon. AGATHA: Automated construction of case law theories through heuristic search. In *Proceedings of the Tenth International Conference on Artificial Intelligence and Law (ICAIL 2005)*, pp. 45–54, New York, NY, USA, 2005. ACM Press.

[18] A. H. Chorley and T. J. M. Bench-Capon. Reasoning with legal cases as theory construction: Some experimental results. In *Proceedings of the Sixteenth Annual Conference on Legal Knowledge and Information Systems (Jurix 2003)*, D. Bourcier, ed., pp. 173–182, Amsterdman, The Netherlands, 2003. IOS Press.

[19] A. H. Chorley and T. J. M. Bench-Capon. AGATHA: Automation of the construction of theories in case law domains. In *Proceedings of the Seventeenth Annual Conference on Legal Knowledge and Information Systems (Jurix 2004)*, pp. 89–98, Amsterdman, The Netherlands, 2004. IOS Press.

[20] C. G. Christie. *The Notion of an Ideal Audience in Legal Argument*. Kluwer Academic Publishers, 2000.

[21] S. Doutre, T. J. M. Bench-Capon, and P. E. Dunne. On the development of argument positions. Technical Report ULCS-05-002, Department of Computer Science, University of Liverpool, 2005.

[22] P. M. 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.

[23] K. Greenwood, T. J. M. Bench-Capon, and P. McBurney. Towards a computational account of persuasion in law. In *Proceedings of the Ninth International Conference on AI and Law (ICAIL 2003)*, pp. 22–31, New York, 2003. ACM Press.

[24] C. Perelman and L. Olbrechts-Tyteca. *The New Rhetoric: A Treatise on Argumentation*. University of Notre Dame Press, IN, 1969.

[25] J. R. Searle. *Rationality in Action*. MIT Press, Cambridge, MA, 2001.

[26] A. Valente. *Legal Knowledge Engineering: A modelling approach*. IOS Press, Amsterdam, 1995.

[27] D. N. Walton. *Argument Schemes for Presumptive Reasoning*. Lawrence Erlbaum Associates, Mahwah, NJ, 1996.

[28] P. H. Winston. *Artificial Intelligence*. Addison Wesley, Reading, MA, 1992.

Received 15 February 2005