

Theory and Practice in AI and Law: A Response to Branting

Katie Atkinson and Trevor Bench-Capon

Department of Computer Science

University of Liverpool

Liverpool L69 7ZF, UK

{katie,tbc}@csc.liv.ac.uk

Abstract.

In this paper we use our previous work which has examined the different levels involved in reasoning about legal cases to examine some challenges to the relevance of current theoretical work in AI and Law made by Branting. In our model we view the process of legal reasoning as being divided into three distinct but interconnected levels of reasoning. These levels involve a bottom layer concerning facts about the world, a top layer concerning legal consequences, and a layer connecting the two, with conclusions at lower levels acting as premises for higher levels. We use our model to explain Branting's observations and show the relation with other strands of work from the AI and Law community.

Keywords. Legal Reasoning, Legal Information Systems, Theory Construction

1. Introduction

One way in which Artificial Intelligence progresses is through a consideration of specific problems. Such problems provide benchmarks which can be used to test existing systems and techniques and to drive new developments. There is, however, a perennial danger in such an approach in that the focus on these problems becomes an end in itself, and the connection with the reality of which the problem is but one example becomes increasingly tenuous. It is therefore necessary from time to time to relate theoretical treatments to the real problems of the domain.

AI and Law is not exempt from such dangers. Over time a distinctive approach to modelling legal reasoning has developed. This strand of work has its origins in the HYPO system of Rissland and Ashley [2], and its developments through the CATO system of Ashley and Aleven [1], and the current work of Ashley and Brüninghaus [3]. This work, firmly grounded in case-based reasoning, has also inspired some of those who come from a rule based and logic background: including Prakken, Sartor, Hage and Bench-Capon. In recent years the approach has begun to include notions of theory construction (originally advocated by McCarty, e.g. [17]), purpose (originally advocated by Berman and Hafner [9]) and dialogue (introduced by Gordon [14]). The approach typically includes the following features:

- Cases are represented as collections of factors. These are not facts, but rather patterns of facts which have some legal significance;
- Legal reasoning is essentially argumentation: on the basis of the factors present in the case, the plaintiff and defendant put forward arguments as to why the case should be decided in their favour;
- These arguments are often presented as a dialectical exchange between the parties;
- A legal decision involves evaluating competing arguments to see which should be accepted;
- Strengths of arguments are based on the purposes served by accepting the argument concerned.

This approach has produced much interesting and sophisticated work, but some doubts remain as to the relevance of this conception to legal reasoning as it is actually reported. These doubts were forcefully raised in a paper by Branting [10] and also formed part of his keynote address at ICAIL 2005. He identified six points which suggested that theory was out of alignment with practice. This is a challenge to which those of us who are working within the approach outlined above need to respond, and this paper is an attempt to do so.

We do so in the context of our recent work on persuasive argumentation in law. We view legal reasoning as taking place on a number of different but interconnected levels, with each level having its own important role to play in the process. In earlier work [6] we have demonstrated this by showing how a general account of practical reasoning can be applied to arguing about legal cases. In that paper we demonstrated our approach by providing a computational reconstruction, using BDI agents, of the reasoning of the majority and dissenting opinions for *Pierson vs Post*, a particular case from property law. The reconstruction of this case suggested that the reasoning involved can be separated into three distinct levels: factual and normative levels and a level connecting the two, with conclusions at one level forming premises at the next, as we discussed in detail [4,6]. In this paper we discuss how this framework can be used to address Branting's challenges, and show how it fits with other accounts from AI and Law to provide some important distinctions in the process of legal reasoning.

In section 2 we discuss Branting's paper [10] and the six points of conflict he identified between the views of the AI and Law research community and the actual practice of law in precedent-based reasoning. In section 3 we give a brief overview of our account of the levels of reasoning involved in legal cases and we discuss how they can help to explain some of the problems identified by Branting. In section 4 we discuss how the observations made in the previous sections fit with the construction and use of theories, as described by Bench-Capon and Sartor in [8], and relate this to some other leading AI and Law systems. Finally, we offer some concluding remarks in section 5.

2. Empirical Investigations Regarding Precedents

In [10] Branting argues that formal and computational models of legal systems do not closely match real world practices. In particular, he maintains that this problem is most prevalent in computational models of precedent-based legal reasoning, as these models do not correspond to the problem-solving practices that exist in the Anglo-American legal system. In order to address this problem, Branting proposes that the AI and Law re-

search community should shift its focus towards task analysis and increased empirical investigation to ensure that computational models of legal practices do in fact reflect real world legal reasoning. In response to this view, we aim to show how our previous investigations into the different levels of reasoning that occur in legal cases can in fact provide a realistic model for computational representation of case-based reasoning and explain the phenomena observed by Branting. Before we discuss how our proposals reflect this view, we first examine in more detail Branting's observations regarding the inconsistencies that exist between research and actual practices involving precedent-based legal reasoning.

2.1. Precedent-Based Reasoning in Practice

In [10] Branting identifies six points of contention that demonstrate how the conduct of practices in courts actually conflicts with widely accepted views of precedent-based reasoning from the research community. We summarise Branting's six observations below.

1. The outcome of disputes on questions of law is highly predictable.

Whereas the current AI and Law approach sees almost all cases as debatable with arguments on both sides, in practice appeals that are brought before courts that operate towards the lower end of the appellate system (which is where most appeals are dealt with) are likely to have predictable outcomes. The predictability of outcomes does decrease the higher the court of jurisdiction, though only a very small number of cases are actually considered at the highest level. Cases considered at the highest level (e.g. the US Supreme Court) are nearly always unique and highly irregular, making their outcome more unpredictable.

2. Uncertainty arises primarily from disagreements about the characterisation of facts, not from ambiguity about the precedential effect of cases.

Whereas the current AI and Law approach takes the factors present as given, and concentrates on the applicability and use of precedents, the role that precedent cases play in appeal decisions is typically clear to Judges. Any uncertainty involved in the decision is primarily manifest through disagreements about the facts of the case and how they should be described, not through the role that the precedents play.

3. Judicial decisions are almost always justified by precedent, not by statutory rules plus principles of construction.

Whereas the current AI and Law approach emphasises interpretation, virtually all appellate decisions are made solely through simple reference to precedents. In cases where there are no precedents then the court may discuss treatises by noted authorities on the matter.

4. Judicial decisions are not frozen dialectic.

Whereas the current AI and Law approach emphasises argument and debate, it is the role of Judges to make decisions about cases that are based upon their inter-

pretation of the law, rather than upon the particular arguments presented to them by attorneys. Decisions are based upon the relevant laws and precedents, not the oral arguments presented.

5. **Cases are not decided upon the basis of teleological factors.**

Whereas some current work in AI and Law looks to justify priorities between arguments through considerations of purpose, judicial decisions rarely consider the underlying social goals and values that are relevant to the decision. Reported decisions are solely based upon consideration of factual evidence and they take no account of teleological factors.

6. **Factor-based argumentation occurs in only a small minority of cases.**

Whereas the current AI and Law approach relies on cases being represented as sets of factors, this is by no means universal. In certain areas of law where there are no well-defined rules (such as tax law) the courts can indeed identify a collection of factors that are relevant to the case. This type of argumentation however, actually occurs in only a small proportion of cases.

Given these observations that seem to reveal inconsistencies between legal reality and the focus of AI and Law research, we believe that we need to shed some light on the reasons as to why these conflicts exist. We do so by turning to our previous work which views legal reasoning as being divided into three levels. In the next section we will briefly recapitulate this work and show how the division of legal reasoning in this manner can account for some of the issues raised by Branting.

3. **Levels of Reasoning in Legal Cases**

In recent work [6] we have explored the representation of legal cases as practical reasoning. We have taken a general account of practical reasoning developed by the authors [5] and applied it to arguing about legal cases. As a concrete example, we provided a reconstruction of the reasoning of the majority and dissenting opinions for a particularly famous legal case in property law: *Pierson vs Post*, 3 Cai R 175 2 Am Dec 264 (Supreme Court of New York, 1805). This reconstruction successfully made use of Belief-Desire-Intention (BDI) agents to replicate the contrasting views involved in the actual decision. An interesting point that emerged from this application was that it showed how the reasoning involved can be separated into three distinct levels: a factual level, a normative level, and a level connecting the two, with conclusions at one level forming premises at the next. In [4] we attempted to generalise these findings by examining the nature and significance of these levels of reasoning in more detail and we provided further examples of their occurrence in law. Here we will give a brief recapitulation of these levels before discussing how they relate to Branting's observations.

3.1. Levels of Reasoning

In the BDI agent reconstruction of *Pierson vs Post* described in [6], the reasoning involved naturally separated into the following three levels:

- **Level 1:** the level of facts about the world, at which desires are derived, which concerns how the law should be.
- **Level 2:** the level at which the legal system connects with the world to achieve these desires, by the appropriate attribution of intermediate concepts.
- **Level 3:** the level of pure legal concepts at which a settled theory is applied.

These levels are connected in that conclusions at lower levels serve as premises at higher levels. Level 3 is concerned solely with legal concepts and the rights they confer. Level 2 is concerned with the ascription of these legal concepts, given the particular facts of the case under consideration. Here arguments for and against ascription of the legal concepts can come either from precedents or from purposes derived from reasoning in the bottom layer (Level 1). At Level 1 people reason about the world in order to determine what the law should be, and conclusions from this level are used at Level 2. These levels and their relationship are shown diagrammatically in Figure 1 below:

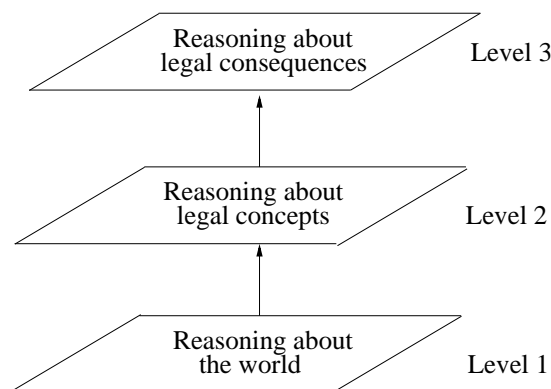


Figure 1. The three levels of legal reasoning emerging from the reconstruction of *Pierson vs Post*.

Related work has also observed these three levels, as we discussed in detail in [4]. In that paper we examined how our levels relate to Lindahl's work [15,16], which makes a distinction between fact situations, intermediate legal concepts, and the consequences that flow from them. We also discussed the work of Breuker and den Haan [12], which explored the need to represent knowledge of the world, knowledge of legal concepts and the connections between them. Additionally we examined Ashley and Brüninghaus' work [3] which stressed the vital role played by intermediate concepts in predicting the

outcome of cases. The idea there is that the decision is predicted on the basis of the intermediate concepts which apply, and the factor based reasoning taken from CATO is mainly used to establish them. This has clear correspondences with our Level 3, where the intermediate concepts are applied, and our Level 2 where the presence of intermediate concepts is established. However, their approach uses case-based reasoning whereas we rely on a model of practical argumentation, and they also make no appeal to purpose. What we add is the ability to motivate the ascription of intermediate concepts through desirable purposes, determined at Level 1.

3.2. *Relation to Branting's Observations*

We now return to Branting's observations regarding the conflicting views of precedent-based reasoning between the research community and the actual practice of law. Here we will describe how each of the points of contention raised by Branting can be explained in terms of the particular levels of reasoning, and the issues involved at each of these levels, that we have identified and described above. A key point, as discussed in [4], is that the decision is very often stated in terms of the reasoning at Level 3, occasionally descending to Level 2 and very rarely (and then typically only at the highest level of appeal) reaching Level 1.

1. **The outcome of disputes on questions of law is highly predictable.**

In our model we can account for this predictability by looking at the reasoning in Level 3. At this level the only reasoning involved concerns the legal decision to be taken, given the facts from the previous levels. Level 3 solely concerns the establishment of a legal conclusion of the case and the issues at this level do not involve disputes over facts. At this stage of the reasoning the facts of the case, including any relevant precedents, have already been established and the decision to be taken should be straightforward, given the relevant applicable laws on the matter. So, we would expect that the outcome of a case would be highly predictable at this level. In *Pierson vs Post*, for example, the majority decision does little more than state that since Post had no possession he has no remedy, which makes the outcome look uncontestable.

2. **Uncertainty arises primarily from disagreements about the characterisation of facts, not from ambiguity about the precedential effect of cases.**

This issue is characterised through the transition from Level 1 to Level 2 reasoning in our model. Here, it must be decided whether the facts of the case in question can instantiate a particular legal concept. The arguments involved are concerned with determining which intermediate concepts apply and also any precedents that are relevant. It is at this stage that the uncertainty of how the facts are to be described must be resolved before being able to proceed to Level 3 reasoning, which solely involves the legal, not factual concepts. Again taking *Pierson vs Post* as our example, the uncertainty on the case revolves around whether Post should be deemed to possess the fox: the facts are not in dispute, only how they

should be characterised in terms of intermediate concepts.

3. Judicial decisions are almost always justified by precedent, not by statutory rules plus principles of construction.

This point relates to Levels 2 and 3 from our model. Here, there are no discussions regarding the purposes and values which motivate the legal decision. Such purposes are considered in Level 1 reasoning, which concerns how the law should be. As the values are fixed once Levels 2 and 3 are reached, the legal decision need make no reference to them, other than to cite the precedent which fixed them.

4. Judicial decisions are not frozen dialectic.

Again, as for the first point, the issue raised here involves the distinction between Level 2 and Level 3 reasoning. As discussed earlier, the reasoning at Level 3 solely involves making legal decisions, given the facts of the case from Level 2, and a settled theory with respect to the precedents. Level 2, where the dialectical exchange of arguments takes place, resolves the questions of interpretation, and only the conclusions are required to state the decision at Level 3.

5. Cases are not decided upon the basis of teleological factors.

As noted in point 3 above, decisions about legal cases do not involve discussion of purposes, social principles and values. In our model, this lack of concern with purpose is reflected by its absence from the reasoning of Levels 2 and 3. However, if we descend to Level 1 reasoning, then this does involve discussion of teleological factors. At this level, facts about the world are discussed and arguments are presented to determine what the law *should* be, if it is to fulfil its purposes. Such discussions typically do not take place in a courtroom (though in the case of *Pierson vs. Post Livingston* (who gave the minority opinion), who needs to extend the existing law to win for his side, does argue a point for what he believes the law should be in terms of purpose). Normally however, such considerations are viewed as a matter to be determined by the legislative body.

6. Factor-based argumentation occurs in only a small minority of cases.

In cases where factor-based argumentation occurs, resolution of disagreements will generally be brought about by descending to Level 1 reasoning. Here, the individual factors are considered and preferences will be formed between the factors, with respect to their individual applicability for the law. As this type of reasoning generally only occurs in areas where there are no well-defined rules, then it is no surprise that the reasoning needs to descend to Level 1. And as Level 1 reasoning is normally taken as settled when a decision is given, then it is also not surprising that such factor-based argumentation only occurs in a small minority of reported decisions. Essentially the factor based considerations are subsumed in the theory being applied, and can be expressed by citing the landmark precedents

in which the debate occurred, without redoing the teleological reasoning which gave rise to the decision.

We believe that the issues raised by Branting can be accounted for in terms of our model identifying the different levels involved in legal reasoning, and which level of reasoning is being reported in the given context. In recognising how and why there are disparities between what we find in AI and Law research and what appears in decisions on cases, we can begin to address the issues in current theories and models of AI and Law.

We now leave our discussion of Branting's work and we turn our attention to the topic of theory construction. In the next section we will show how the levels of reasoning described in section 3.1 are also reflected in the process of theory construction, and also comment on the relation to landmark work in AI and Law, including the HYPO [2] and CATO [1] systems.

4. Relation to Theory Construction

In [8] Bench-Capon and Sartor provide a detailed description of the process involved in constructing theories making use of factors and values. In their model they make use of factors, each associated with some value, which provide a way of describing cases, and these factors are seen as grounding defeasible rules. Preferences between factors are expressed in past decisions, which indicate priorities between the rules. From these priorities it is then possible to abduce certain preferences between the values involved. This model of case law can thus be seen as a process which reveals a preference ordering over values. Bench-Capon and Sartor depicted this process diagrammatically, as shown below in Figure 2. We have added one more box to the original diagram which represents the process of domain analysis required to identify factors and describe cases in terms of them.

As described by Bench-Capon and Sartor, the process of theory construction begins through examination of precedent cases. The next levels are then constructed by identifying the rule-preferences that are revealed in the precedent cases, and the value preferences which these rule preferences show. The theory has now been constructed and it can be used to explain the precedents involved, as well as enabling predictions to be given for new cases.

The levels of legal reasoning we described in section 3.1 fit nicely with Bench-Capon and Sartor's account of theory construction and use in factor-based domains, as we will now elaborate.

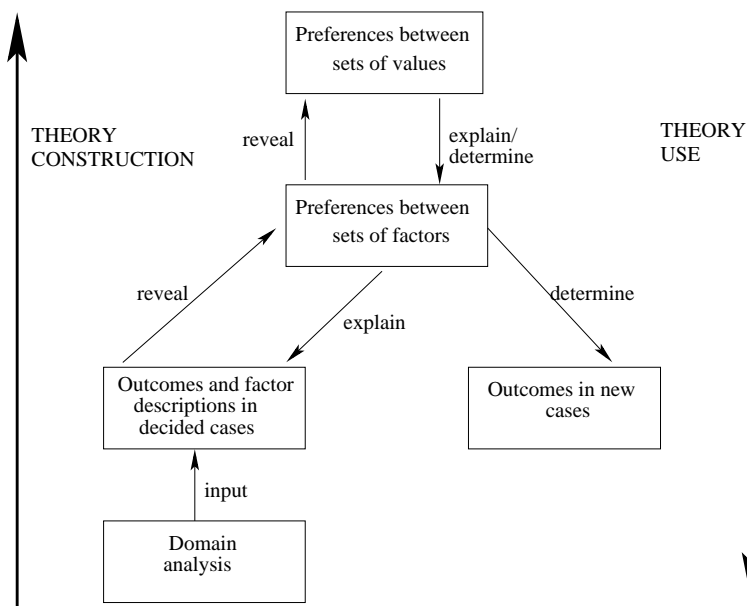


Figure 2. Construction and Use of Theories, adapted from [8].

From Figure 2, we can view the process of factor-based reasoning as an interaction taking place between the different components of the model. At the start, in the bottom left hand corner, Level 1 reasoning is conducted for the domain analysis part of the process. Here, decisions are made to determine which factors are to be included in the theory, which values are served by including them, and which factors are present in the various cases. The relevant factors will be determined by deciding which values are to be promoted and consequently, which factors promote these values. These factors form the input to the next stage, where Level 2 reasoning occurs (the box labelled “Outcomes and factor descriptions in decided cases”). At this point, reasoning is conducted to establish the relative importance and role that each factor played in the decision of the precedent case and this produces a preference ordering on the factors and their combinations. The outcome of this reasoning is represented by the “reveal” relation which takes us from the position where the factors of the precedent cases are considered, to the position where we have a preference ordering between sets of factors (the middle box in the diagram). This preference ordering between the sets of factors will now enable us to establish a value preference ordering, characterised by Level 1 reasoning and shown at the top of the diagram. Recall from the discussions in the previous section that Level 1 is the only level of our model that involves consideration of values. So, the reasoning here reveals the values that are promoted by the factors, though this level of reasoning is generally not of primary concern in the decision making stage of the legal reasoning. Having established a factor and value preference ordering, we now have a complete theory and this can be used to determine outcomes for new cases. This is where the Level 3 reasoning commences and it is at a purely legal level. Here we move down the “determine” relation (on the right hand side of the diagram) where the new legal decisions are made, based upon the constructed theories. If the law is seen as settled, the theory can be taken as given, and the decision in a given case will be couched in terms of the application of this established

theory, and so need contain no reference to values, or to the balancing of factors which led to the ascription of the intermediate concepts.

We can now look at some established systems in AI and Law and see how they address various parts of the process. A simple rule based system, such as [20] is concerned only with the transition from a settled theory to the outcome of cases. Such systems can thus be seen as embodying Level 3 reasoning only, and produce the kind of decisions which confine themselves to the concerns raised by Branting.

A case based system such as CATO involves the transition from precedents to preferences, building arguments as to which intermediate factors (called “abstract factors” in CATO) apply. Ashley’s more recent work with Brüninghaus, the IBP system [3], extends the scope to the application of these preferences, so as also to predict case outcomes, by following the “determines” link to “outcomes in new cases”. CATO goes beyond what is seen in routine decisions, because it is targeted at law students, and needs to address the Level 2 reasoning which explains the ascription of intermediate concepts (represented by CATO’s abstract factors). Argumentation is crucial at this point because the preferences revealed must be tested and justified in terms of their ability to explain and cover the remaining precedents. Prakken and Sartor’s reconstruction of HYPO [19] covers a similar area.

The Level 1 reasoning representing the transition from rule preferences to value preferences introduces teleological notions and was first advocated by Berman and Hafner [9]. These notions are also incorporated in work by Prakken [18] and Bench-Capon [7]. Again the preferences revealed can be disputed and must be justified, accounting for the need for argumentation in this transition. Here, however, we are dealing with the underpinnings of the theory, more the concern of commentaries than judgements, except where we are dealing with a case where the law needs to be extended or changed. Such cases are typically only encountered at the highest levels of appeal. Dialectics are ways of conducting arguments, and so they too will only appear at Levels 1 and 2.

Viewed in this way, we can see the AI and Law approach as attempting to describe legal reasoning in terms of the process required to “reason from first principles”, as opposed to the application of tried and trusted rules. We can see a similar dichotomy in AI approaches to science: where we can either attempt qualitative reasoning on a “deep model”, or use heuristics summarising the behaviour of such a model. In this field attempts have been made to derive the heuristics from the deep model [11]. In AI and Law this can be seen as the underlying approach of [13], where Level 3 and Level 2 reasoning go to construct a theory which is used to produce a Prolog program which is then executed to perform the Level 3 reasoning. As Branting’s first observation indicates, such “first principles” reasoning is deployed in some cases (albeit rare and in the higher courts), so it is legitimate for AI to model this reasoning, even if the fruits of such reasoning can be delivered in less sophisticated systems to handle routine cases.

One transition not addressed in any of these systems is the allocation of factors to cases, represented by the box labelled “domain analysis” in Figure 2. This is germane to Branting’s second criticism, that uncertainty often arise from disputes about the factual characterisation of cases. CATO-type systems do address the attribution of intermediate concepts on the basis of factors, but the factors are taken as given. This may well point to a significant gap in the current theoretical work, since it is the mapping from “raw” facts to these stylised fact patterns that is a key part of the lawyers skill. The use of dimensions in HYPO, allows some limited consideration of this problem, in that its dimensions allow

for some debate as to whether a factor is present (e.g. posing the question: just what is required before we accept that the plaintiff took reasonable measures to protect his “Trade Secret?”). None the less we are still a long way from techniques allowing us to move from the kind of narrative that a client might present, to the kind of fact description that a lawyer can derive from such narratives, as described in detail in [10].

5. Concluding Remarks

In this paper we have used our previous work which has examined the different levels involved in reasoning about legal cases [4]. The presence of these different levels emerged from a BDI agent reconstruction of the famous property law case of *Pierson vs Post*. However, these levels have also proved to be instructive when considering their relation to other research from the AI and Law community. We examined some observations made by Branting in [10] which reveal apparent inconsistencies between legal practice and computational models and theories of AI and Law. We believe that our model can reveal where and why these inconsistencies arise and thus identify important distinctions that need to be made when modelling legal reasoning. We have also used this account to explain some differences in scope in some of the existing leading systems in AI and Law.

Acknowledgements

Katie Atkinson is grateful for support from the EPSRC. Trevor Bench-Capon is grateful for partial support from the European Commission, through project APSIC (IST-FP6-002307).

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 D. Bourcier, editor, *Proceedings of the Sixteenth Annual Conference on Legal Knowledge and Information Systems (Jurix 2003)*, pages 153–162, Amsterdam, The Netherlands, 2003. IOS Press.
- [4] K. Atkinson and T. Bench-Capon. Levels of reasoning with legal cases. In P. E. Dunne and T. Bench-Capon, editors, *ICAAIL 2005 Workshop on Argumentation in AI and Law*, IAAIL Workshop Series, pages 1–11, Nijmegen, The Netherlands, 2005. Wolf Legal Publishers.
- [5] K. Atkinson, T. Bench-Capon, and P. McBurney. Justifying practical reasoning. In F. Grasso, C. Reed, and G. Carenini, editors, *Proceedings of the Fourth International Workshop on Computational Models of Natural Argument (CMNA 2004)*, pages 87–90, Valencia, Spain, 2004.
- [6] K. Atkinson, T. Bench-Capon, and P. McBurney. Arguing about cases as practical reasoning. In *Proceedings of the Tenth International Conference on Artificial Intelligence and Law (ICAAIL 2005)*, pages 35–44, New York, NY, USA, 2005. ACM Press.
- [7] T. Bench-Capon. Try to see it my way: Modelling persuasion in legal discourse. *Artificial Intelligence and Law*, 11 (4):271–87, 2003.

- [8] T. Bench-Capon and G. Sartor. A model of legal reasoning with cases incorporating theories and values. *Artificial Intelligence*, 150 1-2:97–143, 2003.
- [9] 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)*, pages 50–59, New York, NY, USA, 1993. ACM Press.
- [10] K. L. Branting. An agenda for empirical research in AI and law. In *Working Papers of the ICAIL 2003 Workshop on Evaluation of Legal Reasoning and Problem-Solving Systems*, pages 28–35, Edinburgh, UK, 2003.
- [11] I. Bratko, I. Mozetič, and N. Lavrač. *KARDIO: a study in deep and qualitative knowledge for expert systems*. MIT Press, Cambridge, MA, USA, 1989.
- [12] J. Breuker and N. den Haan. Separating world and regulation knowledge, where is the logic? In *Proceedings of the Tenth International Conference on Artificial Intelligence and Law (ICAIL 1991)*, pages 92–97. ACM Press: New York, USA, 1991.
- [13] A. 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)*, pages 45–54, New York, NY, USA, 2005. ACM Press.
- [14] T. F. Gordon. The Pleadings Game: An exercise in computational dialectics. *Artificial Intelligence and Law*, 2:239–292, 1994.
- [15] L. Lindahl. Deduction and justification in the law. The role of legal terms and concepts. *Ratio Juris*, 17 2:182–202, 2004.
- [16] L. Lindahl and J. Odelstad. Normative positions within an algebraic approach to normative systems. *Journal of Applied Logic*, 17 2(1):63–91, 2004.
- [17] L. T. McCarty. An implementation of *Eisner v Macomber*. In *Proceedings of the Fifth International Conference on AI and Law (ICAIL 1995)*, pages 276–286, New York, NY, USA, 1995. ACM Press.
- [18] H. Prakken. An exercise in formalising teleological case-based reasoning. *Artificial Intelligence and Law*, 10 (1-3):113–133, 2002.
- [19] H. Prakken and G. Sartor. Modelling reasoning with precedents in a formal dialogue game. *Artificial Intelligence and Law*, 6 (2-4):231–287, 1998.
- [20] M. J. Sergot, F. Sadri, R. A. Kowalski, F. Kriwaczek, P. Hammond, and H. T. Cory. The British Nationality Act as a logic program. *Communications of the ACM*, 29(5):370–386, 1986.