Argumentation in AI and Law:

what do we know and where should we go?

Workshop held in conjunction with the Fourteenth International Conference on Artificial Intelligence and Law

Casa dell'Aviatore,
Viale dell'Università 20,
Rome

Monday June 10th, 2013

Chair: Trevor Bench-Capon, University of Liverpool

Programme

9.00-9.15: Introduction to the Workshop: Trevor Bench-Capon

Session 1: Tools for Legal Argumentation

1a: Tools for Legal Argumentation: What do we know?

9.15- 9.45 Tools for Legal Argumentation: Current State of the Art: Speaker: *Thomas F. Gordon*, *Fraunhofer Institute for Open Communications Systems (FOKUS), Berlin*
9.45-10.00 Commentary: *Henry Prakken, Univerisities of Utrecht and Groningen*
10.00-10.45 Discussion
10.45-11.00 Summary

11.00-11.30 Coffee Break

1.b: Tools for Legal Argumentation: Where should we go?

11.30-12.15: Panel Session: Tools for Legal Argumentation: What Issues should we be Addressing?
Chair: *Vern R Walker, Hofstra University*
Panellists:
Chris Reed, *University of Dundee*;
Anne Gardner, *Independent Scholar*;
Marc Lauritsen, *Capstone Practice Systems*.
12.15-12.45 Open Discussion: Tools for Legal Argumentation: A Research Agenda
12.45-13.00 Summary

13.00-14.00 Lunch

Session 2: Legal Argumentation Using Cases

2a: Legal Argumentation Using Cases: What do we know?

14.00- 14.30 Legal Argumentation Using Cases: Current State of the Art. Speaker: *Trevor Bench-Capon, University of Liverpool*
14.30-14.45 Commentary: *Kevin Ashley, University of Pittsburgh*
14.45-15.30 Discussion
15.30-15.45 Summary

15.45-16.15 Coffee Break

2.b: Legal Argumentation Using Cases: Where should we go?

16.15-17.00: Panel Session: Legal Argumentation Using Cases: What Issues Should We Be Addressing?
Chair: *Edwina Rissland, University of Amherst*
Panellists:
Matthias Grabmair, *University of Pittsburgh*;
Padmaja Sasidharan, *King's College, London*;
Josh Blackman, *South Texas College of Law*

17.00-17.30 Open Discussion: Legal Argumentation Using Cases: A Research Agenda
17.30-17.45 Summary

17.45-18.00 Workshop Close

CONTENTS

Tools for Legal Argumentation: Current State of the Art

Tom Gordon

Fraunhofer FOKUS /
University of Potsdam

Factual and Legal Issues

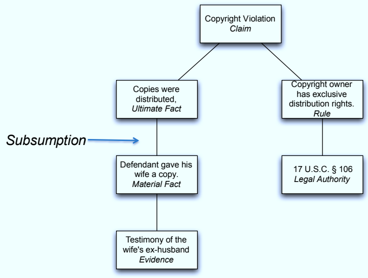

Legal Argumentation Tasks

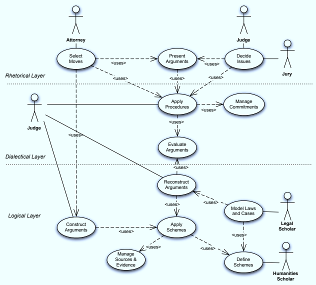

Bank XX (1993/96)
Rissland, Skalak & Friedman

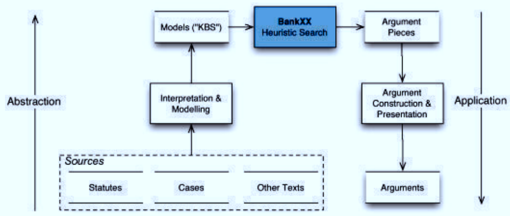

PLAID (1995)
Bench-Capon & Staniford

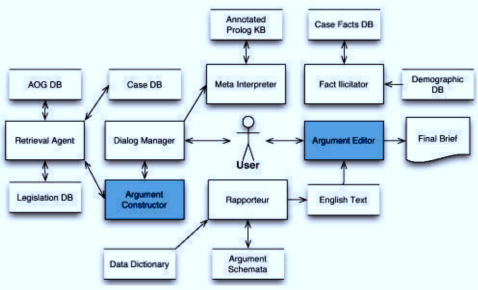

Legal Argumentation Tasks

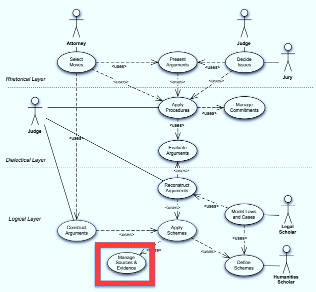

Managing Sources and Evidence

- Legal Research Services
  - Google Scholar
  - LexisNexis
  - WestLaw
- Content, Knowledge and Case Management Systems
  - Alfresco
  - Drupal
  - Plone
- Markup and Metadata
  - CEN MetaLex
  - Akoma Ntoso
  - OASIS Legal Document XML

Legal Argumentation Tasks

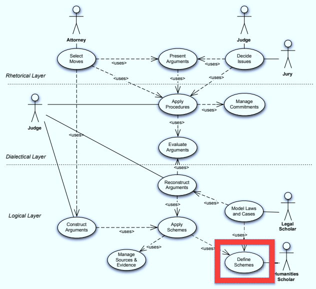

Computational Models of Argumentation Schemes

- Argument from Cases (CBR) [McCarty, Ashley, Rissland, Branting, Skalak, Aleven, Roth]
- Argument from Rules and Cases [Gardner, Branting, Skalak, Nitta, Prakken, Sartor, Bench-Capon]
- Argument from Rules with Priorities [Hage, Verheij, Gordon, Prakken, Sartor]
- Argument from Rationales [Loui, Norman, Roth]
- Argument from Principles [Hage, Prakken, Gordon, Loui]
- Argument from Values, Purpose and Policy [Berman, Hafner, Bench-Capon, Sartor]
- Argument from Goals [Atkinson, Bench-Capon]
- Argument from Evidence [Prakken, Walton]

Argumentation Scheme Languages

Araucaria

```
<SCHEME>
<NAME> Argument from Expert Opinion</NAME>
<FORM>
<PREMISE>E is an expert in domain D</PREMISE>
<PREMISE>E asserts that A is known to be true</PREMISE>
<PREMISE>A is within D</PREMISE>
<CONCLUSION>A may (plausibly) be taken to be true.</CONCLUSION>
</FORM>
<CQ>Is E a genuine expert in D?</CQ>
<CQ>Did E really assert that A is known to be true?</CQ>
<CQ>Is the expert's pronouncement directly quoted? If not, is a reference
to the original source given? Can it be checked?
</CQ>
<CQ>If the expert advice is not quoted, does it look like important
information or qualifications may have been left out?</CQ>
<CQ>If more than one expert source has been cited, is each authority
quoted separately? Could there be disagreements among the cited
authorities?</CQ>
<CQ>Is what the authority said clear? Are there technical terms used that
are not explained clearly? If the advice is in layman's terms, could this be
an indication that it has been translated from some other form of
expression given by the expert?</CQ>
<CQ>Is A relevant to domain D?</CQ>
<CQ>Is A consistent with what other experts in D say?</CQ>
<CQ>Is A consistent with known evidence in D?</CQ>
</SCHEME>
```

Carneades

```
(make-scheme
  :id 'expert-opinion
  :header (make-metadata
        :title "Argument from Expert Opinion"
        :source "Douglas Walton, Legal Argumentation and
              Evidence, The Pennsylvania State University Press, University Park,
              2002, pp.49-50.  Douglas Walton, Appeal to Expert Opinion, The
              Pennsylvania University Press, University Park, Albany, 1997,
              p.211-225.")
  :conclusion '?A
  :premises [(make-premise
                 :role "major"
                 :statement '(expert ?E ?S))
             (make-premise
                 :role "domain"
                 :statement '(in-domain ?A ?S))
            (make-premise
                 :role "minor"
                 :statement '(asserts ?E ?A))]
   ;; Critical Questions
  :exceptions [(make-premise
                    :role "CQ1"
                    :statement '(untrustworthy ?E))
               (make-premise
                   :role "CQ2"
                  :statement '(inconsistent-with-other-experts ?A))]
 :assumptions [(make-premise
                   :role "CQ3"
                   :statement '(based-on-evidence ?A))])
```

Legal Argumentation Tasks

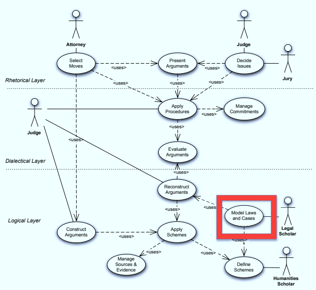

Modelling Laws

- Isomorphism
- Reification
- Defeasibility
- Contraposition
- Case-Based Reasoning
- Rule Validity
- Modalities

Some Rule Languages for Modeling Laws

- Defeasible Logic (Nute 1994; Governatori, Rotolo & Sartor 2005)
- PRATOR (Prakken & Sartor 1996)
- Legal Knowledge Interchange Format (Gordon et al. 2008)
- OASIS Legal RuleML (2013 ?)
- Carneades Scheme Language (Gordon 2013)

Example Rule

Legal RuleML

```
<ruleml:Rule key=":rule1" closure="universal">
   <lrml:hasStrength> <lrml:Defeasible/> </lrml:hasStrength>
   <ruleml:if>
    <ruleml:Atom>
      <ruleml:Rel iri="#person"/> <ruleml:Var>X</ruleml:Var>
    </ruleml:Atom>
   </ruleml:if>
   <ruleml:then>
      <lrml:SuborderList>
         <lrml:Prohibition>
            <ruleml:Atom>
               <ruleml:Rel iri="#engageCreditActivity"/>
               <ruleml:Var>X</ruleml:Var>
            </ruleml:Atom>
         </lrml:Prohibition>
      </lrml:SuborderList>
   </ruleml:then>
</ruleml:Rule>
```

Carneades Scheme Language

```
(make-scheme
   :id 'rule1
   :strict false
   :pro true
   :conclusion '(prohibited (engage-credit-acitivity ?X))
   :premises [(pm '(person ?X))])
```

Modeling Cases

- Title
- Court
- Issue
- Decision
- Facts or Factors
- Arguments (majority and minority)
- Ratio Decidendi

HYPO Trade Secrets Example

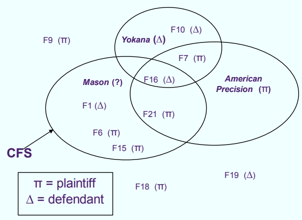

CASE Yokana (Δ)
- **F7 Brought-Tools (π)**
- **F10 Secrets-Disclosed-Outsiders (Δ)**
- **F16 Info-Reverse-Engineerable (Δ)**

CASE American Precision (π)
- **F7 Brought-Tools (π)**
- **F16 Info-Reverse-Engineerable (Δ)**
- **F21 Knew-Info-Confidential (π)**

CASE Mason (CFS, Undecided)
- **F1 Disclosure-in-Negotiations (Δ)**
- **F6 Security-Measures (π)**
- **F15 Unique-Product (π)**
- **F16 Info-Reverse-Engineerable (Δ)**
- **F21 Knew-Info-Confidential (π**

Modelling Arguments of Cases Example: Popov v. Hayashi

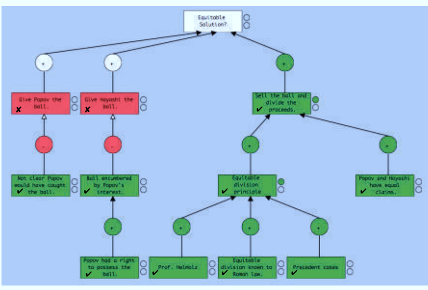

Gordon, T. F., and Walton, D. A Carneades Reconstruction of Popov v Hayashi. *Artificial Intelligence and Law 20*, 1 (2012), 37–56.

Ratio Decidendi: Theory Construction

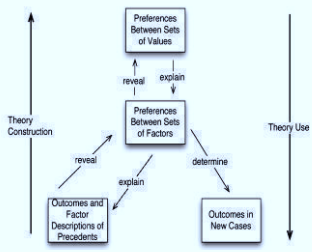

Bench-Capon, T., and Sartor, G. A Model of Legal Reasoning with Cases Incorporating Theories and Values. *Artificial Intelligence 150*, 1–2 (Nov. 2003), 97–143.

Legal Argumentation Tasks

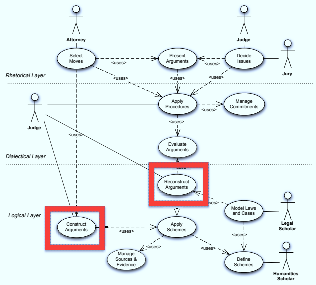

Constructing and Reconstructing Arguments

- Compared
  - Construction: creating original arguments by instantiating argumentation schemes
  - Reconstruction: using argument schemes to interpret existing arguments in natural language texts (e.g. court opinions)
- Kinds of Tools
  - Interactive software tools
  - Fully automatic, using models of, e.g., facts, ontologies, rules and cases

Interactive Argument Reconstruction with Araucaria

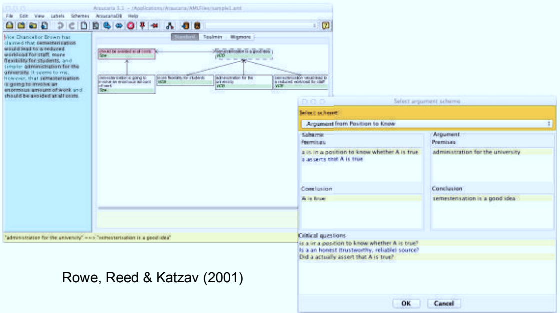

Rowe, Reed & Katzav (2001)

Argument Mining: Automatic Argument Reconstruction

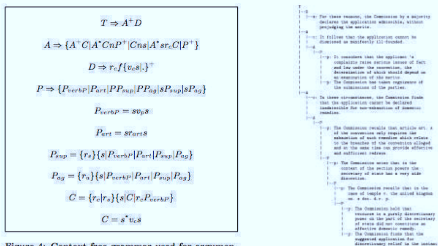

Figure 4: Context-free grammar used for argumentation structure detection and proposition classification

Figure 5: Tree Structure of an argument

Palau, Raquel Mochales, and Marie-Francine Moens. "Argumentation mining: the detection, classification and structure of arguments in text." Proceedings of the 12th international conference on artificial intelligence and law. ACM, 2009.

Automatic Argument Construction from Rules and Ontologies

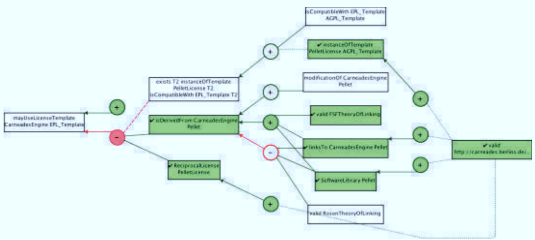

Gordon, T. F. Combining Rules and Ontologies with Carneades. In *Proceedings of the 5th International RuleML2011@BRF Challenge* (2011), pp. 103–110.

Legal Argumentation Tasks

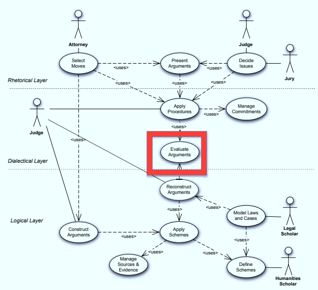

Evaluating Arguments: Conceptions of Argument

- Single-step arguments: Instantiations of argumentation schemes
- Defeasible proofs (Pollock 1987; Prakken 2010)
- Minor premise (Pragma-Dialectics)
- Set of propositions (Bresnard & Hunter 2008)
- Argument graphs (Gordon, Prakken & Walton 2007)

Evaluating Arguments: Procedure

1. Validate that each single-step argument properly instantiates its scheme. Check for missing premises.
2. From the perspective of the **audience** of interest, such as a judge or jury, label the statements which are accepted as true, or rejected as false, without argument, and weigh/order the single-step arguments.
3. **Narrower conception of evaluation**: Evaluate the defeasible proofs in the argument graph to determine which arguments are acceptable (in), not acceptable (out) or undecided. Use this information to then compute, analogously, which of the statements (claims) are acceptable (in), not acceptable (out) or undecided.
4. Use argumentation schemes to reveal and critically question any implicit premises and to construct counterarguments.

Computational Models of Argument Evaluation

- Narrow conception of evaluation
- Abstract Arguments
  - Abstract Argumentation Frameworks (Dung 1995)
  - Value-based Argumentation (Bench-Capon 2003)
  - Using arguments about preferences (Modgil 2009)
- Structured Arguments
  - DefLog (Verheij, 2003)
  - Using proof standards; Carneades (Gordon, Prakken & Walton 2007)
  - Defeasible proof trees; ASPIC+ (Prakken 2010)
  - Mapping Carneades to ASPIC+ (Gizjel & Prakken 2011)

ArguMed 3 (2001) Verheij

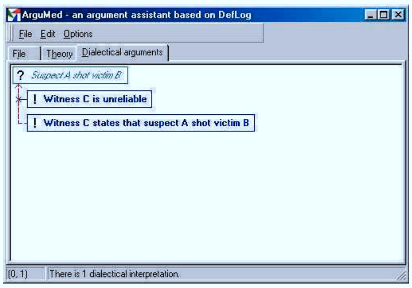

Carneades 2011

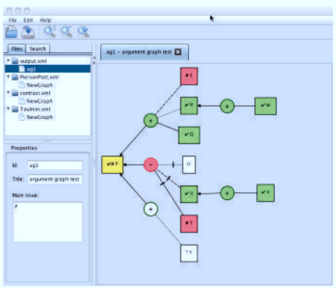

TOAST

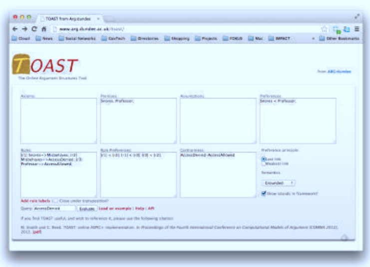

M. Snaith and C. Reed. TOAST: online ASPIC+ implementation. In Proceedings of the Fourth International Conference on Computational Models of Argument (COMMA 2012), 2012.

TOAST

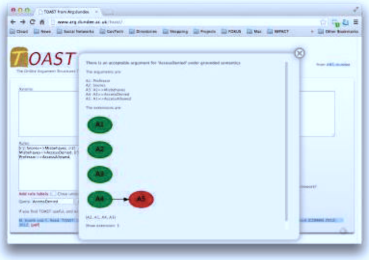

Legal Argumentation Tasks

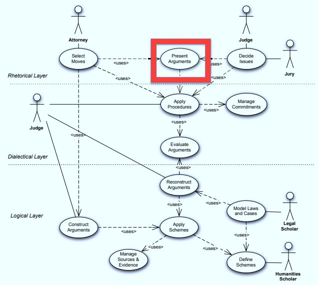

Presenting Arguments

- Textually
  - Outlines
  - Hypertext
  - Reports, using “document assembly” tools (e.g. HotDocs, Exari)
- Diagrams
  - Argument maps

Hypertext Outline
Carneades Web App (2012)

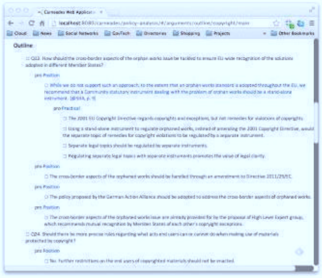

Detailed Argument View
Carneades Web App

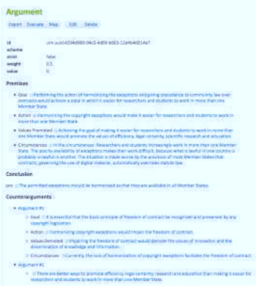

Rationale (2003)
Austhink

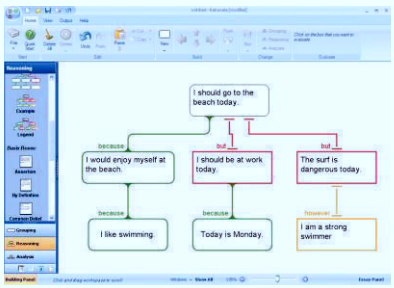

Carneades Argument Map
Web App Version

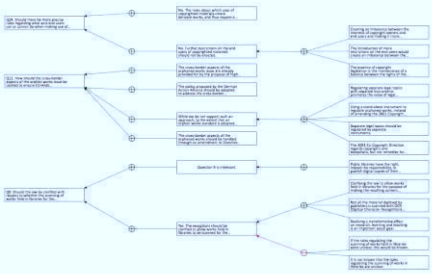

ArguNet (Betz)

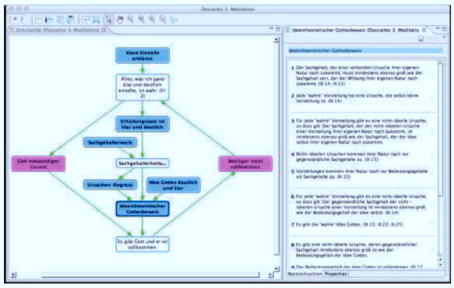

LASAD

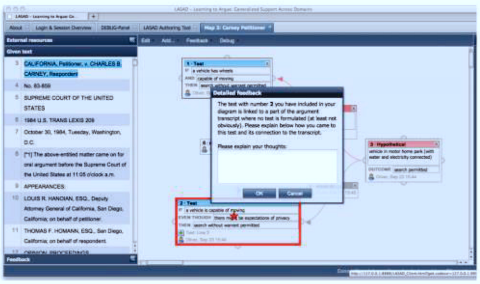

Loll, Frank, and Niels Pinkwart. "LASAD: Flexible representations for computer-based collaborative argumentation." International Journal of Human-Computer Studies (2012).

Tasks Not Covered

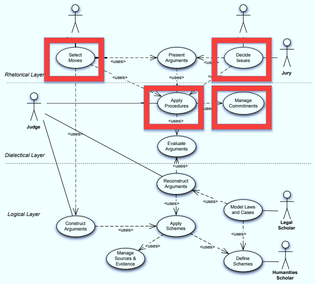

Whither Argumentation Tools? I have my suspicions.

Chris Reed
University of Dundee
Scotland

There is a busy, dynamic, driven academic community producing theory, tools and techniques of argumentation which by and large are having no impact on the world. Why not? Let's examine the SUSPECT.

- **Scale**. There is a general challenge for work transitioning from the academic lab to the real world: where in the lab a technique may work for a toy example of ten or twenty parts, it needs to work for ten or twenty thousand; ten or twenty million. We need to face up to the fact that this scaling up is not just a minor engineering problem -- in some cases it may be as big a problem as the one we initially set out to solve.

- **Users**. It is encouraging to see a few examples of people working with users to develop tools they actually want -- but it's still rare. Working with practising lawyers, for example, can give insights that can't come from the lab alone.

- **Seductive**. The tool must be one that users want to use and go on using..

- **Pushmepullyou**. In a similar vein, there may be philosophically pleasing theory, or mathematically elegant proofs, but for delivering tools it has to be a case of (at least some) end-user pull, not just technology push. We have to be solving the problems people have, not the problems we'd like them to have.

- **Entrepreneurship**. To really make tool deployment fly we need committed entrepreneurs -- not necessarily people in it for the money (witness the Debategraph team, for example), and not necessarily people outside the academic sphere (withness the OU group, for example), but people who can commit to getting things in to the hands of users.

- **Crowd**. The last two are more about tech trends. The first is that silo-ed argumentation -- a single user using some tool as they might use Word -- is on the wane. Argumentation in all its guises is increasingly a digital social activity. I don't mean it's all about arguing on facebook: sharing rulings with other law makers; sharing the decisions of cases with other case-deciders; sharing procedure with other followers of procedure: this social interaction is a core part of what the law is about, and our tools should reflect that.

- **Text**. Lawyers and the law, perhaps even as much as academics, are driven by text. Our tools, though, are typically boxes and arrows, propositions and such. As argumentative text processing starts to climb on the academic agenda, it will be exciting to see tools in this space starting to become available.

Rules and Cases for Legal Reasoning: Notes on Some Neglected Aspects

Anne v.d.L. Gardner

286 Selby Lane, Atherton, California 94027
gardner@cs.stanford.edu

Abstract

Rules and cases are essential elements in legal reasoning, but computational models have barely begun to reflect the complexities of their roles. Based on experience with a real case, this paper identifies four areas that deserve attention from anyone concerned with understanding the processes of a general legal reasoner. These are (1) combining rules that were adopted for differing purposes but that all have application to the problem at hand; (2) allowing for argument over the logical structure of rules, and managing to reason with them even when unsure what the logical structure is; (3) allowing cases to be used mainly for their facts and outcome, mainly for their reasoning, or mainly for the rules they lay down, and employing each technique when appropriate; and (4) extending the legal sources that are treated as cases. The paper does not propose solutions but merely attempts, by way of examples, to suggest significant research areas.

Introduction

It is widely agreed that a realistic computational model of legal reasoning must use both rules and cases. Working out how to combine them has been the subject of a number of programs, including GREBE (Branting 1991a, 1991b, 1994), CABARET (Rissland and Skalak 1991; Skalak and Rissland 1992), TAXMAN II (McCarty and Sridharan 1981; McCarty 1989, 1995), and my own dissertation project (Gardner 1987).

The choice of directions for developing this work is goal-dependent. One may want mainly to build tools useful for lawyers; to understand human cognitive processing; or to understand the forms of legal analysis, argument, and decision-making. Toward the last goal at least, it is important to observe what moves take place in actual legal reasoning. Our models so far have mostly been based on abstractions—inspired, for instance, by the descriptions in standard works of jurisprudence like Levi (1949) and Hart (1961). According to my recent experience with a large, rather technical, real-life case, the proportion of the reasoning that existing computational models can account for is disappointingly small. This is not only for the expected reasons, such as inability to handle general natural language and commonsense knowledge (on the AI side) or arguments from purpose and from principle (on the legal side). There are also less familiar features, some fundamental. This paper identifies a few features, of varying importance, that are fairly easy to pick out and describe.

Background: The *Alaska* Case

United States v. Alaska was a lawsuit over the ownership of lands just off Alaska's north coast. The areas in dispute are potentially valuable for oil; the government that owns them is the government that gets to decide whether to open the lands for offshore oil exploration and, if opened, to enter leases with oil companies and collect royalties from them.

Being a suit between a state and the federal government, the *Alaska* case was initially filed in the Supreme Court as an original jurisdiction case under Article III, sec. 2, clause 2, of the Constitution. The Court, as it often does in such cases, appointed a Special Master to hear the case and report back to it. The Master's report (565 pages, covering six main groups of questions) was submitted in March 1996. The Supreme Court, after hearing oral argument by the parties on the parts of the report to which they took exception, issued its decision in June 1997 (117 S. Ct. 1888). Alaska's three exceptions were overruled; the United States' one exception was sustained. My work in all this was with the Special Master, Professor J. Keith Mann of the Stanford Law School.

Combining Rule Sets

Programs that work with statutes usually deal with the terms of a single enactment, for example the British Nationality Act (Sergot et al. 1986) or the Latent Damage Act (Susskind 1989). When cases are used as well, the rule-based part may be limited to a single statutory section, such as the Internal Revenue Code section on home office deductions (Rissland and Skalak 1991).

In *United States v. Alaska*, the range of rules was much broader. The basic statute was the 1953 Submerged Lands Act (43 U.S.C. §§ 1301–1315), which says that each state owns the submerged lands in a three-mile belt measured outward from its coastline. Secondary was a 1958 treaty, the Convention on the Territorial Sea and the Contiguous Zone (15 U.S.T. 1606), which defines the baselines from which nations are to measure their territorial seas. Linking the two was a Supreme Court decision saying, roughly, that

AAAI Spring Symposium Series, 1998:
Multimodal Reasoning

"coastline" in the Submerged Lands Act should be interpreted to mean the same thing as "baseline" in the Convention. (*United States v. California*, 381 U.S. 139 (1965).) So far so good. It is not hard to imagine a rule base that includes both the Submerged Lands Act and the Convention. A bit more detail is given in Gardner (1989).

But of the four main sections of the Master's report that applied these rules—each to a different geographical feature that might or might not form part of the coastline—only two were able to use just the rules mentioned and the related cases. The other two sections involved possible exceptions to the rules, that is, arguments that the usual definitions should not be applied in this particular situation. These arguments were based on other rules entirely.

For example, one part of the case involved a mile-long pier built out into the ocean. Is the three-mile limit to be measured from the mainland or from the end of the pier? Under the Convention, the Master's report found, the end of the pier should be used. The United States, arguing for an exception, invoked sources including the statute authorizing construction of the Trans-Alaska pipeline (43 U.S.C. §§ 1651–1655), the statute giving the Army Corps of Engineers authority over structures built in navigable waters (33 U.S.C. § 403), regulations issued by the Army Corps of Engineers (33 C.F.R. § 209.120 (1976)), and some judge-made rules about what happens when government employees fail to follow government regulations.

This example is not extraordinary. Altogether the Master's report contains citations to over forty different statutes. It could not have been known in advance just what statutes should form part of the rule base.

The question raised here is not just how to make a broader rule base available to programs. What reasoning is involved in figuring out how disparate rule sets fit together? Is it true that introducing a new rule set is usually associated with arguing for an exception to the main rule? If so, does this tell us anything new about defeasible reasoning?

Interpretation of Rules

There are some important operations in legal reasoning that are not covered by the general concepts of applying rules and analogizing cases. One is the interpretation of a rule, in the sense of reasoning about the meaning of a universally quantified proposition. Such reasoning may needed as a step separate from reasoning about whether the antecedents of a rule are satisfied by the particulars of a case at hand. One problem of rule interpretation, familiar from the work of Allen and Saxon (e.g., 1987, 1991), arises from ambiguity in the natural-language counterparts of logical operators.

The need for rule interpretation goes further, however; it may involve clarifying the relations among domain concepts. Some examples come from the Convention's definition of a bay. (Where a bay is found, a line drawn across its mouth counts as part of the coastline.) The definition contains two sentences, with the first stating some general conditions and the second imposing a geometrical test based on the area of a semicircle:

> *Article 7(2)*. For the purpose of these articles, a bay is a well-marked indentation whose penetration is in such proportion to the width of its mouth as to contain landlocked waters and constitute more than a mere curvature of the coast. An indentation shall not, however, be regarded as a bay unless its area is as large as, or larger than, that of the semi-circle whose diameter is a line drawn across the mouth of that indentation.

The interpretation problems do not come from the rule-plus-exception structure of the English, which can uncontroversially be flattened to

if general-conditions(x) and semicircle-test(x)
then bay(x) .

Rather, the problems in the *Alaska* case were (1) does the semicircle test subsume the general conditions? and (2) if not, what is the logical structure of the general conditions? Dependent on the answers was the ownership of lands under a water body called southern Harrison Bay. It was agreed that southern Harrison Bay met the semicircle test.

Alaska's first argument was, in effect, that

if semicircle-test(x) then general-conditions(x) ,

thus reducing the rule to

if semicircle-test(x) then bay(x) .

This is a startling position because it has no warrant in the syntax of the English rule; but it is not preposterous, because scholars have raised the same question of interpretation (see Report, pp. 182–83). In Alaska's view, the interpretation was warranted by the drafting history of the definition. The Master's report, after reviewing the history, disagreed.

The lesson here is that even where a rule has an authoritative text with no surface structural ambiguity, programs still need to leave room for argument over what logical expression correctly translates the rule. Had the Master's report found that the drafting history supported Alaska's argument, there would then have been a need for metalevel reasoning about whether preferring the history to the syntax is legitimate (for some legal sources, see Report, p. 186, n. 11). The latter point goes beyond the suggestion in Rissland and Skalak (1991) that one may always argue for dropping an antecedent from a rule.

Once it is decided that *general-conditions(SoHarBay)* must be tested, the second set of problems arises. Part of the difficulty comes from the usual source, namely vague or open-textured predicates such as "well-marked indentation," "penetration," "landlocked waters," and "more than a mere curvature of the coast." (Another predicate, "the width of its mouth," presented no problem in this case because the value was agreed to be about 12 nautical miles.) But beyond this, the relationship among the predicates is unclear: how do we write the rule

if $p_1(x)$ *and . . .* $p_n(x)$ *then general-conditions(x)* ?

From the English, one might think that "well-marked indentation" and "more than a mere curvature of the coast" were separate requirements. As used in the drafting history, however, they seem synonymous; and indeed the general conditions were criticized as circular during the drafting (Report, p. 191). In addition, the relationship between "penetration" and "landlocked waters" is unclear. Are they independent requirements, or does the value of one determine the value of the other? Supreme Court precedents seem to lead to the first conclusion; the syntax, to the other (see Report, pp. 199–200).

In the Master's report, the outcome of this second exercise in rule interpretation was in effect to conclude that we do not know the logical form of the first English sentence in the definition of a bay. This conclusion triggered a case-based approach to testing the general conditions, using the various predicates as factors rather than as a neat conjunction of preconditions.

Modes of Reasoning with Cases

Programs that work with legal cases usually apply the same algorithm no matter what the source of the case. Human reasoners, however, adapt their style to the situation. If the court hearing the current case has issued some recent decisions that are more or less on point, an extremely elaborate analysis may be called for, spelling out every point of similarity and dissimilarity and attending to every nuance in what the court said. In other situations one may use a precedent only for its facts and its result, ignoring the reasoning. In still others the important thing may be the pattern of reasoning displayed in the precedent rather than close factual similarity. The Master's report contains examples of all of these. Here are illustrations of the latter two.

Fact-based Comparison

Continuing with the example of southern Harrison Bay, the most difficult feature to reach a conclusion about was whether the waters were landlocked. As shown in figure 1, the area has two arms, where both parties agreed the waters were landlocked, and a middle area, on which they disagreed. The most relevant precedents were other indentations with two arms, sometimes called double-headed bays. Five such precedents were available (covering one formation in Norway, one in California, and three in Alaska including the northwestern part of Harrison Bay). Given maps of each precedent area, together with the decision on whether it formed a single bay, it was possible to order the precedents along a numerical scale (Report, pp. 216–26). Southern Harrison Bay fell within the landlocked range; the Master's report recommended accordingly; and the parties filed no exceptions.

This portion of the report seems unusual in making so little use of the reasoning in past cases and in producing a

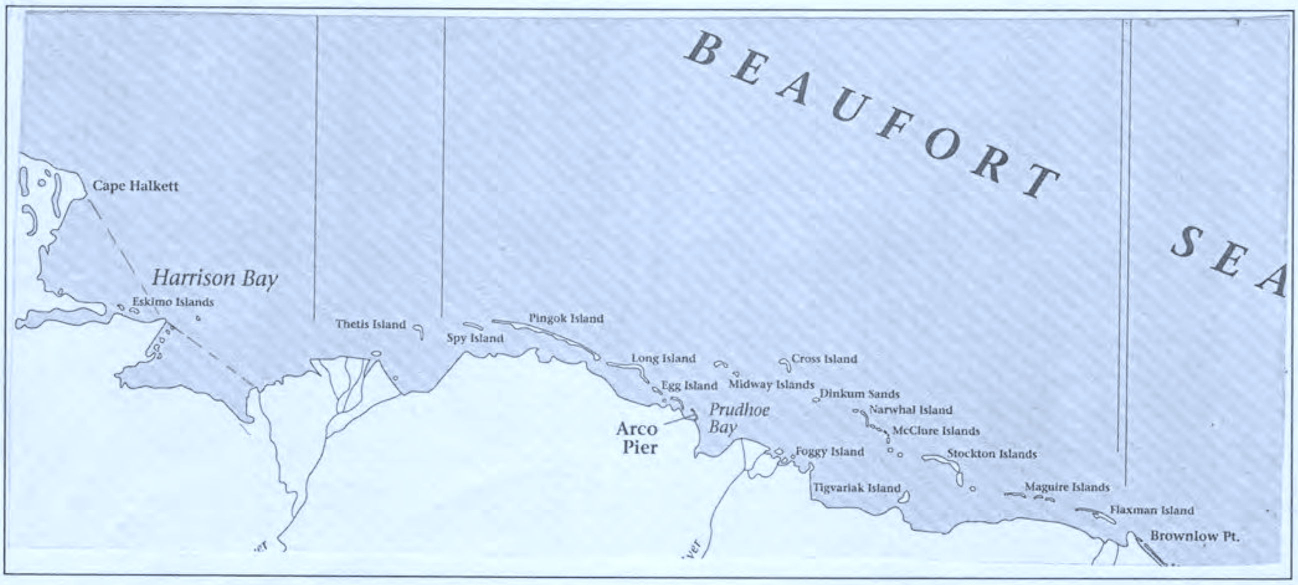

Figure 1. A section of the north coast of Alaska. Disputed areas include the southern part of Harrison Bay (east of the Eskimo Islands), the islands around Prudhoe Bay, and the Arco pier. The northwestern part of Harrison Bay was agreed to form a bay.

basically geometrical solution to a legal problem. There are several explanatory factors. First, the precedents on double-headed bays did not contain much usable reasoning: only one was the subject of a judicial opinion, and even that one, from the International Court of Justice, did little more than announce its conclusion that the Norwegian Svaerholthavet had "the character of a bay" (see Report, p. 207).

Such guidance as there was came instead from a United States Supreme Court decision in which the geographical facts had less resemblance to southern Harrison Bay. In that decision (on whether Long Island Sound and the adjacent Block Island Sound formed a bay), the Court identified some requirements for landlocked waters:

> We agree with the general proposition that the term "landlocked" "implies both that there shall be land in all but one direction and also that it should be close enough at all points to provide [a seaman] with shelter from all but that one direction." [*Rhode Island and New York Boundary Case*, 469 U.S. 504, 525 (1985)]

In the report, an elaboration of "land in all but one direction" led to the numerical scale used to decide whether southern Harrison Bay was landlocked. The question whether the land was "close enough at all points" was essentially answered by article 7(4) of the Convention, which permits a line across the mouth of a bay to be as much as 24 miles long. As a final simplification, the evidence on the extent of shelter for a mariner was limited to the information available from two-dimensional maps. The last move seems to be standard, both in the legal cases and also, no doubt for practical reasons, among geographers who need to draw boundaries.

Reason-based Comparison

An interesting contrast to the Harrison Bay reasoning comes from a different part of the case, this one on the effect that near-shore islands have on drawing the coastline (see figure 1). The United States said that each island has its own coastline and its own three-mile limit. Alaska wanted to draw a single line as the coastline, running along the seaward side of any islands, and measure the three-mile limit from there. For these results the United States invoked the Convention, while Alaska argued on various grounds that the usual rules of the Convention should not apply.

One suggestion by Alaska was that the islands should be treated in the same way as islands along the most similar parts of the United States coast, notably in Louisiana and Mississippi. This would have been in line with the reasoning used for Harrison Bay. As the analysis was worked out in the Master's report, however, the exact geography in past cases proved less important than the theory behind each decision on how to treat an island configuration. The theories that were used were sometimes unclear, mutually inconsistent, and unpredictable even in application to the same area at different times. Consequently there was no basis for inferring how the islands off the north coast of Alaska would have been treated before the Convention took effect, and thus no basis for applying an exception to the Convention's rules. This was so even though a 1985 Supreme Court case contained a statement that strongly supported Alaska's position:

> Prior to its ratification of the Convention on March 24, 1961, the United States had adopted a policy of enclosing as inland waters those areas between the mainland and off-lying islands that were so closely grouped that no entrance exceeded 10 geographical miles. [*Alabama and Mississippi Boundary Case*, 470 U.S. 93, 106 (1985)]

The Master's report (pp. 53–54) took the statement to be nonbinding because (1) it was a statement of fact, not of law, and (2) it was not strictly necessary to the 1985 decision. The Supreme Court agreed and overruled Alaska's exception.

The example highlights the importance of working on representations for the reasoning in legal cases, not just for the facts and the outcome. As for the differences from the Harrison Bay example—with respect to the role of the key sentence from the most important precedent, and with respect to the treatment of geographically similar cases—these are not inconsistencies. Rather, they result from differences in the available evidence, the available cases, and the arguments based on these that were or could have been made.

Sources of Cases

For a human reasoner, one of the most satisfying moments is finding a case that solves a puzzle or clinches an argument. Case-based programs may have the same goal. Achieving the goal, however, often requires going beyond the case base that seems natural for the problem at hand. Perhaps it is precisely because the normal stock of judicial opinions provides no answer that an issue becomes salient as a puzzle. Examples from the Master's report of nonstandard cases include the following:

1. To help settle the meaning of "permanent" in the Convention, a case from a domain having nothing to do with submerged lands but holding that eight years is long enough to count as permanent. (Report, p. 320.) The case was cited in one of the parties' briefs; the brief writer might have found it from the legal reference *Words and Phrases*.

2. To help interpret an early boundary description, a case that had matching facts but that turned on another point. The statement relevant to our problem was thus dictum, and moreover it appeared in a concurring opinion.

(Report, p. 378.) The case was located through an early Supreme Court opinion (*Shively v. Bowlby*, 152 U.S. 1 (1894)), which provided a virtual treatise on the submerged lands decisions to that time.

3. To help establish the meaning of "high tide" in the Convention, a decree entered in a previous case. There had been no discussion of the question in the opinions leading up to the decree, but the decree itself equated "high tide" with "mean high water." (Report, p. 234.) A Supreme Court decree also yielded one of the examples of a double-headed bay, and others provided some of the information on past treatment of islands as part of the coastline or not. The decrees are published in United States Reports, the same source as for Supreme Court opinions.

4. To help settle whether a body of water qualified as a bay under the Convention, the minutes of a meeting of a committee of federal officials, deciding that a similar neighboring body of water was a bay. (Report, p. 225.) The minutes were a document introduced into evidence—as were, again, many of the items pertaining to past treatment of islands. This suggests, at least for a case in which much of the evidence is documentary, that the line between precedent cases and the facts of a current case is less sharp than usually supposed: building the case base may require processing some of the evidence.

Conclusion

In this paper I have tried to identify some significant differences between human legal reasoning and the computational models we have so far. If the paper is successful, it will have suggested some fruitful directions for future research.

References

Allen, L. E., and Saxon, C. S. 1987. Some Problems in Designing Expert Systems To Aid Legal Reasoning. In Proceedings, First International Conference on Artificial Intelligence and Law, Boston, 94–103. New York: Association for Computing Machinery.

Allen, L. E., and Saxon, C. S. 1991. More IA Needed in AI: Interpretation Assistance for Coping with the Problem of Multiple Structural Interpretations. In Proceedings, Third International Conference on Artificial Intelligence and Law, Oxford, 53–61. New York: Association for Computing Machinery.

Branting, L. K. 1991a. Integrating Rules and Precedents for Classification and Explanation: Automating Legal Analysis. Ph.D. diss., University of Texas at Austin.

Branting, L. K. 1991b. Reasoning with Portions of Precedents. In Proceedings, Third International Conference on Artificial Intelligence and Law, Oxford, 145–154. New York: Association for Computing Machinery.

Branting, L. K. 1994. A Computational Model of Ratio Decidendi. *Artificial Intelligence and Law* 2: 1–31.

Gardner, A. v.d.L. 1987. *An Artificial Intelligence Approach to Legal Reasoning*. Cambridge: MIT Press.

Gardner, A. v.d.L. 1989. Representing Developing Legal Doctrine. In Proceedings, Second International Conference on Artificial Intelligence and Law, Vancouver, 16–22. New York: Association for Computing Machinery. Also in A. Martino, ed., *Expert Systems in Law*, 201–209. Elsevier Science Publishers, 1992.

Hart, H. L. A. 1961. *The Concept of Law*. Oxford: Clarendon Press.

Levi, E. H. 1949. *An Introduction to Legal Reasoning*. Chicago: University of Chicago Press.

McCarty, L. T. 1989. A Language for Legal Discourse: I. Basic Features. In Proceedings, Second International Conference on Artificial Intelligence and Law, Vancouver, 180–188. New York: Association for Computing Machinery.

McCarty, L. T. 1995. An Implementation of Eisner v. Macomber. In Proceedings, Fifth International Conference on Artificial Intelligence and Law, College Park, Maryland, 276–286. New York: Association for Computing Machinery.

McCarty, L. T., and Sridharan, N. S. 1981. The Representation of an Evolving System of Legal Concepts: II. Prototypes and Deformations. In Proceedings, Seventh International Joint Conference on Artificial Intelligence, Vancouver, 246–253.

Rissland, E. L., and Skalak, D. B. 1991. CABARET: Rule Interpretation in a Hybrid Architecture. *International Journal of Man-Machine Studies* 34: 839–887.

Sergot, M. J.; Sadri, F.; Kowalski, R. A.; Kriwaczek, F.; Hammond, P.; and Cory, H. T. 1986. The British Nationality Act as a Logic Program. *Communications of the ACM* 29: 370–386.

Skalak, D. B., and Rissland, E. L. 1992. Arguments and Cases: An Inevitable Intertwining. *Artificial Intelligence and Law* 1: 3–44.

Susskind, R. E. 1989. The Latent Damage System: A Jurisprudential Analysis. In Proceedings, Second International Conference on Artificial Intelligence and Law, Vancouver, 23–32. New York: Association for Computing Machinery.

United States v. Alaska (No. 84, Original), Report of Special Master J. Keith Mann (1996), *exceptions overruled in part and sustained in part,* 117 S. Ct. 1888 (1997). A limited number of copies of the report are available from Anne Gardner. Major law libraries may have the report either in hard copy or on microfiche.

Words and Phrases. 1940–. St. Paul, Minn.: West Publishing Co. 46 vols.

Arguing about Balances

A position paper for the **Workshop on Argumentation and AI and Law**
International Conference on AI and Law, Rome, Italy, June 2013

Marc Lauritsen
Capstone Practice Systems
Harvard, Massachusetts USA

Many practical arguments end up being about which position has the strongest set of good reasons in its favor. In other words, which way an imaginary balance tips once all the considerations pro and con the various arguments are stacked onto corresponding pans of a metaphorical scale. Rather than embodying a complex structure of attack and support relationships, with rebutters, undercutters, and underminers, some arguments simply come down to debates about how much 'weight' the considerations deserve and which outcomes are most favored once all have been taken into account. Rather than being most naturally modeled in inference trees, such arguments lend themselves to representation in simple arrays, in which reasons for competing outcomes are summed and compared.

The same is often true in contexts of public deliberation, where candidate policies are judged against multiple criteria and where constituents differ both in terms of the relative 'goodness' of candidates on the criteria and the relative importance of the criteria. To use an example explored by Bench-Capon, Prakken, and Visser at ICAIL 2011 ('Argument schemes for two-phase democratic deliberation'), imagine two citizens with different views of two possible strategies for reducing deaths:

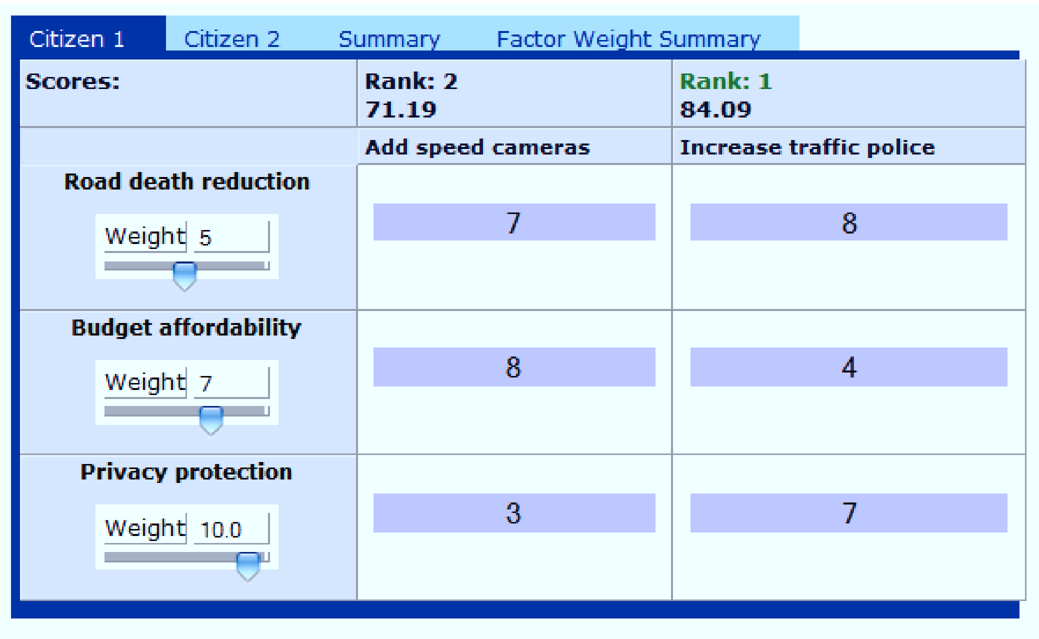

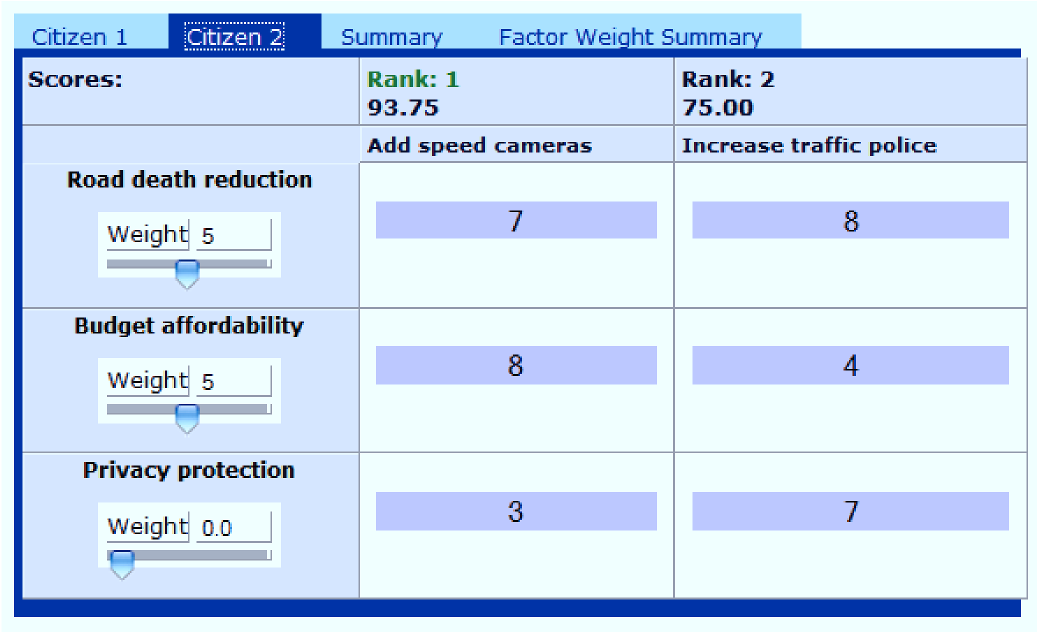

Citizen 1	Citizen 2	Summary	Factor Weight Summary

Scores:	Rank: 1 93.75	Rank: 2 75.00
	Add speed cameras	Increase traffic police
Road death reduction Weight 5	7	8
Budget affordability Weight 5	8	4
Privacy protection Weight 0.0	3	7

These two citizens happen to agree on their assessment of the *effect* of the policies, but differ on the *values* that are implicated. You can imagine them expressing their views more graphically, where widths and heights are used rather than numbers for effect assessments and value weights respectively, and total ‘volumes of relative goodness’ are shown at top:

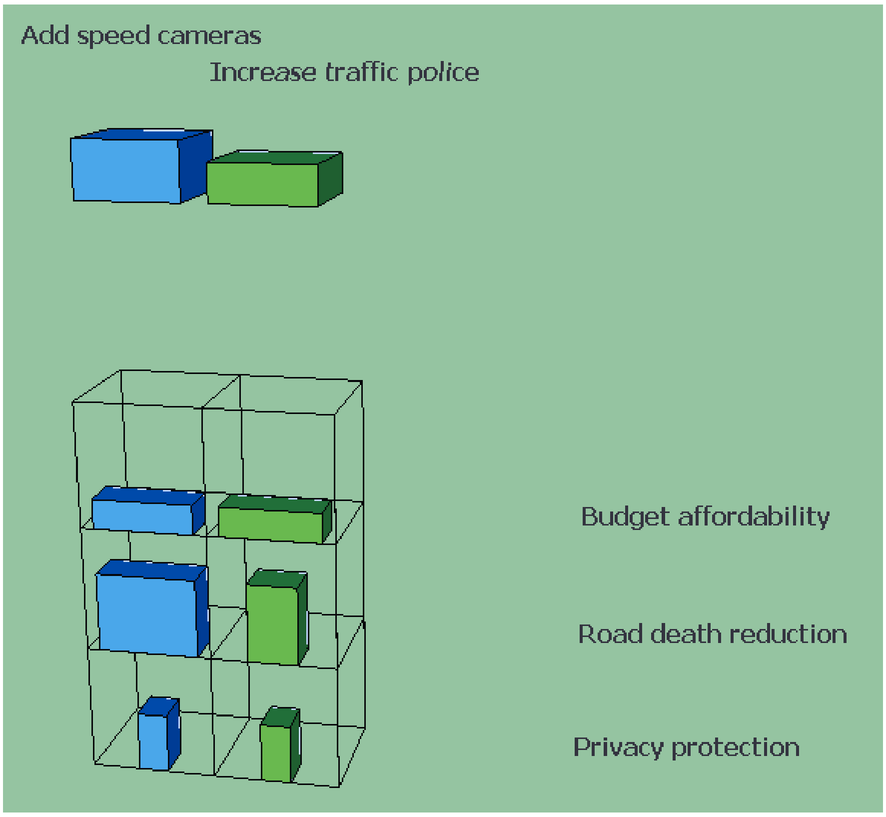

The author has pursued this approach to modeling arguments and deliberations about balances in his work on a ‘choiceboxing’ system. The attached case study summarizes how that system supported recent deliberations by a national community of experts on the uses of technology in legal services for the poor in the United States. In that process thirty people expressed views about ten possible initiatives, judged against eleven criteria. The case study provides a concise summary of the method and an encouraging example of its successful deployment.

For purposes of this workshop, the follow questions are suggested:

1. In addition to its apparent practical uses, does choiceboxing facilitate theoretical investigations? Might it serve e.g. as a convenient way to express at least the states in an argument or deliberation ‘game’? Can such games be understood as states of a choicebox and transitions between them?

2. Are many arguments in legal contexts adequately modeled by choicebox-style representations?

3. What aspects of argumentation and deliberation go beyond what can effectively be reflected in such models? (Note that the full choiceboxing system includes value functions that accommodate non-linear utility curves and necessary/sufficient logic.)

4. How might preferred semantics and other formal frameworks be used to impart greater intelligence to environments in which participants manipulate choicebox-like representations of their claims and positions as they argue and deliberate?

Arguing with Legal Cases: what do we know?


Trevor Bench-Capon
Department of Computer Science
University of Liverpool
UK

Landmarks – Part I

- ***Taxman***: Thorne McCarty
  - Theory Construction, Prototypes and Deformations
    - *Eisner v Macomber* (Tax Law Case)
- ***HYPO***: Edwina Rissland and Kevin Ashley
  - Three Ply Argument, Dimensions
    - US Trade Secret Law

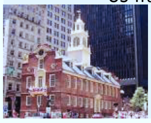

Landmarks – Part 2

- ***CABARET***: Edwina Rissland and David Skalak
  - Top Level of Rules, Argument Moves
    - Home Office Deduction
- ***CATO***: Kevin Ashley and Vincent Aleven
  - Factors, Abstract Factors, Down-Playing
    - US Trade Secret Law
- ICAIL 1993: Don Berman and Carole Hafner
  - Purpose, Teleological Reasoning
    - Pierson v Post, Keeble v Hickersgill, Young v Hitchens
      - *The Wild Animals Cases*

Landmarks – Part 3

- ***Rule Based Representation of Precedent:*** Henry Prakken and Giovanni Sartor
  - **Precedents represented as three rules:** plaintiff rule, defendant rule and priority between them
    - **Residence example (Fictional)**
- ***Value Based Theories*:** Trevor Bench-Capon and Giovanni Sartor and Henry Prakken
  - **Rule Preferences explained as value preferences, Theory Constructors**
    - **Wild Animals Cases**

Landmarks – Part 4

- ***IBP***: Stefanie Bruninghaus and Kevin Ashley
  - Top Level of Issues: Prediction
    - US Trade Secrets Law
- ***Argumentation Schemes***: Adam Wyner, Katie Atkinson, Trevor Bench-Capon, Henry Prakken
  - Reasoning with Cases as Practical Reasoning
  - Reconstruction of CATO using Argumentation Schemes
    - Wild Animals + *Popov v Hayashi*

Some Other Approaches – Part 1

- ***GREBE***: Karl Branting
  - Semantic Networks
    - Industrial Injury
- ***Neural Networks***:
  - Daniele Bourcier
  - Trevor Bench-Capon
    - Hospital Visit Benefit (fictional)

Some Other Approaches – Part 2

- ***BANKXX:*** Edwina Rissland, David Skalak and M. Timur Friedman
  - Assembling Arguments through Heuristic Search
    - US Bankruptcy Law
- ***AGATHA***: Alison Chorley and Trevor Bench-Capon
  - Constructing Theories as an Adversarial Game
    - Wild Animals, US Trade Secrets

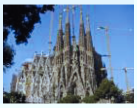

Some Other Approaches – Part 3

- ***Tangled Hierarchies:*** Bram Roth and Bart Verheij
  - Attacks on connections as well as factors
    - Dutch Employment Law
- ***Evidence***: Floris Bex, Peter van Koppen, Susan van den Braak, Henry Prakken and BartVerheij
  - Resolving conflicting witness testimony
    - Criminal cases

The Problem

- Given a set of decided cases, and a new (current) case: ***how we construct and compare arguments about how the new case?***
- **Some Issues:**
  - ***Where do we start****?*: facts (higher courts) or evidence (lower courts)?
  - ***How do we represent cases?*** Facts, Dimensions, Factors, Issues, Purposes, Values.
  - ***How do we compare cases?***
  - ***How do we go beyond* a fortiori *reasoning?***

From Evidence to Decision

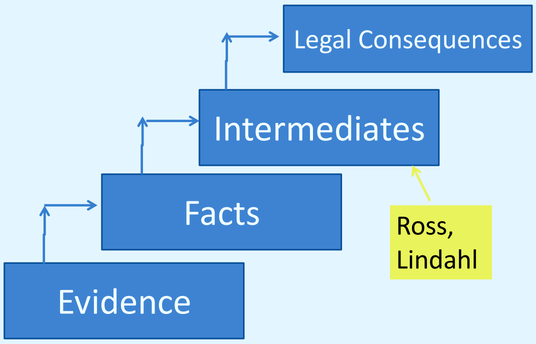

What makes cases similar?

- ***Not*** closeness of fact
  - Woman of 60
  - Woman of 59
    - Not similar
  - Woman of 60
  - Man of 65
    - Similar
  - Woman of 62
  - Man of 62
    - Not Similar
  - Woman of 98
  - Man of 67
    - Similar

So cannot use standard techniques like least squares

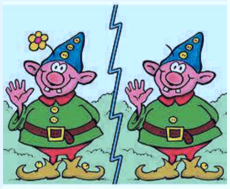

Why is a raven like a writing desk?

- That is from Lewis Carrol's Alice in Wonderland, but there are real cases in which match:
  - A fox
  - A shoal of fish
  - A baseball
- ***All being pursued***
  - A lemonade bottle
  - A coffee urn
  - A car with a loose wheel
- All imminently dangerous

Facts are unique to their cases

We need abstractions which are ***legally relevant***

Factors provide such abstractions

CATO

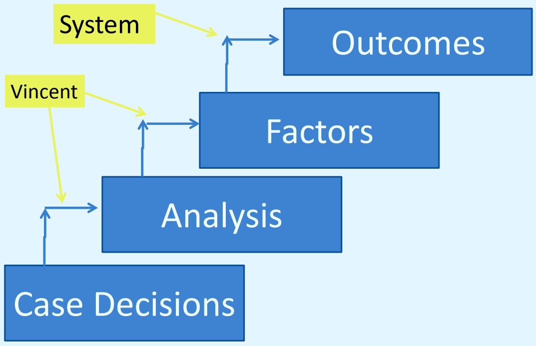

Factors

- Legally relevant features of cases: abstract from many, disparate, fact patterns
  - Emerge from case law
    - Judges relate to previous decisions through similar language: these are the factors
  - Favour one side or the other
    - Unlike dimensions
  - Are determined by analysis
    - Attempts to automate largely unsuccessful (e.g. SMILE, Bruninghaus and Ashley)

Factor Based Reasoning

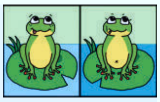

- Cases are represented as sets of factors
- One step of inference:
  - Antecedent is a conjunction of factors
  - Consequent is an outcome
- Exact matches are rare: precedents are ***distinguished*** when
  - a precedent is cited for the one side
  - the precedent is stronger for that side
  - the current case is weaker for that side

Factors May Result In

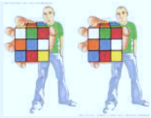

- An *a fortiori* argument for one side
  - Precedent which cannot be distinguished: all opposing precedents can be distinguished
    - (a,b,c: ?): (a,b: plaintiff) : (a,b,d: defendant)
- Arguments for both sides
  - No *a fortiori* precedent for either side
  - Distinguishable precedents for one or both sides

How do we extend our theory to choose between competing arguments?

Beyond *A Fortiori*

- CATO – Abstract Factor Hierarchy
  - Factors are children of more abstract factors
  - Factors may be reasons ***for*** ***or against*** the presence of their parent
    - ***NOT an IS-A hierachy***
  - Children factors may ***substitute*** for or ***cancel*** one another (*downplaying*)

Value Based Theory Construction

- Factors are associated with social ***values*** (purposes)
  - Deciding for the part favoured by the factor would promote the value
  - Precedents express preferences between values
  - These value preferences form a theory which explains the past decisions (more or less well)
    - We choose the best theory
  - The theory is applied to the current case
  - Can also be ***independent*** arguments for value preferences

Horty and Bench-Capon (2012)

- In Prakken and Sartor (1998) both the plaintiff and the defendant rules were as strong as possible (used ***all*** the factors)
- In Horty and Bench-Capon (2012) the rule for the winning side may be weaker (use only a subset of the factors)
  - This means that it can apply to more cases
  - But can only be justified by success

Organising Factors

- Often cases are seen as sets of factors. But often too there is some organisation into topics or issues.
  - **CABARET**: Top level logical expression representing the statute rule
    - **Factors interpret the terms of the statute**
  - **IBP**: Top level logical "model" (from Restatement of Torts)
    - **Factors are partitioned into issues to resolve the terms of the model**
  - **Theory Construction:** Factors relate to values
    - **Factors determine which values can be promoted: preferences decide which values will be promoted**

Organising Arguments

- Three Ply Argument (HYPO, CATO)
  - Citation
  - Distinguishing and Counter Examples
  - Rebuttal: Distinguishing Counter Examples etc.
- Dialectic Tree (e.g. Prakken and Sartor)
  - Argument for
  - Argument against
  - And so on
- Cascade of Argumentation Schemes (e.g. Wyner et al)
  - Top Level Scheme
  - Schemes to establish premises of higher schemes
  - Schemes to undercut higher schemes

What About Dimensions?

- Factors as Points on Dimensions:

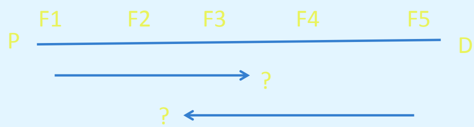

Where does the cross over point come?

Where do factors come from? Which side do they favour?

Other Roles for Dimensions?

- Perhaps Dimensions connect to
  - Abstract Factors?
  - Issues?
  - Values?
  - Elements in Tests?

How do Factors Combine?

- Using Logical Connectives?
  - Top Level provides Necessary and Sufficient Conditions
  - Top Level Provides Argument – some elements may be missing
- "Considerations "
  - The factors need to be ***weighed*** against one another and a judgement made

Similar considerations apply to Values

And how do we compare sets of Values?

Summary

- We understand reasoning from factors to outcomes reasonably well
  - Why we need factors
  - The logic of precedent
  - Where factors fit in the overall process
- We have some understanding of what we need to investigate
- We have some ideas about how to go about these investigations
- We have no clarity or consensus on these areas

Towards Measurable Intelligent Inference

Position paper on the future of legal argumentation with cases in AI&Law

Matthias Grabmair
Intelligent Systems Program
University of Pittsburgh, USA

Introduction

Research on modeling legal argumentation with cases has explored various ways to represent cases and arguments about them. Strong connections to computational models of argument exist with regard to representation and inference/semantics. While many insights have been gained, I will argue in this short paper that, from the perspective of a potential future user of a legal-case-argumentation tool, the state of the art in the field does not yet enable the development of systems capable of drawing intelligent and useful inferences from available knowledge. In order to tackle this challenge, I advocate for a more detailed exploration of value-based/purposive reasoning and the corresponding knowledge representation problems as well as for a commitment to an implementation and empirical evaluation of developed formalisms.

Limitations of Current Knowledge and Value Representations

Formal models of legal argumentation with cases have been construed with the assumption of having certain formalized knowledge available. The most prominent representation of cases is that of dimensions/factors as originally introduced by HYPO [3] and taken further by CATO [1] and IBP [4] as well as by the theory construction model [9]. Factors/dimensions are stereotypical fact patterns in the domain of discourse and of potential relevance to the case and need to be manually encoded into the representation. However, even a significant advancement in this natural-language-processing problem of recognizing archetypical fact patterns in case descriptions would not by itself move the available formalisms into a position of making a practical contribution. This is because formal models of legal arguments with cases have not sufficiently explored how lawyers argue about *why* and *how* the presence or absence of certain facts in a case affect the decision. This standard of ¨sufficiency¨ shall be understood as the capacity of the models and formalisms we develop to interface with related technologies to construct an intelligent application for the purpose of increasing the productivity of its user in a task related to arguing with legal cases. Aside from abundant work in legal theory and methodology, AI&Law researchers have made considerable efforts to include teleology into formal models of legal argumentation. The necessity for a system¨s capacity to reason substantively about teleology has been explained initially by Berman & Hafner [10] and has since been tackled in various ways. CATO introduced a hand crafted factor hierarchy to generate more complex arguments. IBP grouped factors into issues. Factors can also be associated with values [12] to make their presence and absence from cases more informative. Theory construction uses an abstract ordering of values to prioritize rules with which new cases can be decided. Such an ordering can also steer inference in value-based argumentation frameworks [6]. Recently, values have been further examined in the practical reasoning setting [5] and with regard to rule-based argumentation with thresholds [8]. Also, our work on the value judgment formalism [11] uses argumentation with values and effects on them to enhance the representation of purposive reasoning about the impact of facts and legal rules.

Beyond Factors: Exploring the Building Blocks of Value-Based Reasoning

Recent AI&Law work on argumentation with values [7] has established connections between formal models and US Supreme Court Jurisprudence (i.e. legal practice) as well as legal theory work [13], respectively. AI&Law¨s contribution, however, must be significantly more granular and practical than existing legal theory formalisms (such as, e.g., Alexy¨s work on argumentation with cases [2]) and at the same time enable practitioners to use legal expert systems to their benefit. In other words, the next generation of work on legal argumentation with cases should ideally be both computational (as opposed to purely representational) and suitable for an empirical evaluation of its achievements. I see the next step towards such a contribution in decomposing argumentation with values into its functional elements at a greater level of granularity than current representations. Designing factor-based systems (or comparable knowledge representations) involves significant domain expertise in both encoding the factors as well as interrelating them in a meaningful way to allow for the desired level of teleological reasoning capability. It appears to be a more promising goal to explore, implement and evaluate formalisms based on patterns of value-based reasoning across legal domains and strive to assemble a vocabulary compatible with that of general purpose knowledge representations or semantic extraction from natural language. For example, in our most recent work on the value judgment formalism [11], we model argumentation about the relevance of fact patterns in cases by identifying their effects in the domain of discourse and connecting them to values. We thereby describe the relevance of legal concepts in a more general vocabulary and open up the representation towards general causal, temporal and agent-based reasoning as well as a contextual balancing of values. The needed knowledge is still immense, but may be more modular as well as easier to maintain and extend. Conceptually similar recent work exists in value-based argumentative practical reasoning [5].

Empirical Evaluation of Inference Capacity

While the correctness of predicting case outcomes can be evaluated in a straightforward way, there are no established evaluation metrics for the generation of intelligent arguments in AI&Law. The lack of data corpuses exacerbates the difficulty of developing suitable systems and conducting informative experiments. However, it should be a fruitful endeavor to discuss which inference capacity the field is striving for. Which kinds of systems do we want to build? How would their inference look like? How would a prototype be evaluated? At the same time, it will be worth to work on specifying the precise needs that a knowledge base would need to fulfill in order to allow productive system development. What kinds of ontologies would be needed? What else is necessary? Does it need to be perfect or is some degree of error tolerable? Exploring these questions may provide guidance for further developments, allow for interchange with other areas of AI and help gradually introduce a notion of empirical validity into the field.

Conclusions

In this paper, I have argued that research on modeling legal argumentation with cases in AI&Law has reached the point where a significant advancement towards a practical contribution is best facilitated if (1) representation of and argumentation with values and purposes are explored in greater detail, and (2) if these efforts are guided by a commitment to implementation and empirical evaluation. I look forward to seeing more work focusing on developing and evaluating more fine-grained representations of cases, facts, values and their interaction so that we can move closer towards systems capable of autonomously generating intelligent legal arguments with cases.

References

[1] Aleven, V., Using Background Knowledge in Case-Based Legal Reasoning, Art. Int. 150(1-2), 183-238, 2003.
[2] Alexy, R., Two or Three?, in: Martin Borowski (ed.), On the Nature of Legal Principles, Archives for Philosophy of Law and Social Philosophy, Beiheft 119, pp. 9-18, Steiner, 2010.
[3] Ashley, K.D., Ontological Requirements for analogical, teleological and hypothetical reasoning, ICAIL 2009 Proceedings, pp. 1-10. ACM Press, 2009.
[4] Ashley, K. & Brüninghaus, S., A Predictive Role for Intermediate Legal Concepts, Proc. Jurix 2003, pp.153-162, IOS Press 2003.
[5] Atkinson, K., Bench-Capon, T., Cartwright, D. and Wyner, A., Semantic models for policy deliberation, ICAIL 2011 Proceedings, pp.81-90, ACM, 2011.
[6] Bench-Capon, T., Persuasion in Practical Argument Using Value-based Argumentation Frameworks, J. Logic & Comp., 13(3):429-448, 2003.
[7] Bench-Capon, T., Relating values in a Series of Supreme Court Decisions. In K. Atkinson (editor): Legal Knowledge and Information Systems. JURIX 2011 Proceedings, pp. 31-22, IOS Press 2011.
[8] Bench-Capon, T. & Prakken, H., Using argument schemes for hypothetical reasoning in law, Artificial Intelligence and Law, Vol. 18, Issue 2, pp. 153-174, Springer, 2010.
[9] Bench-Capon, T. & Sartor, G., A model of legal reasoning with cases incorporating theories and values, Artificial Intelligence, Vol. 150, Issue 1-2, pp. 97-143, Elsevier, 2003.
[10] Berman, D.H. and Hafner, C.D., Representing teleological structure in case-based legal reasoning: the missing link, ICAIL 1993 Proceedings, pp. 50-59, ACM, 1993.
[11] Grabmair, M. & Ashley, K.D., Using Event Progression to Enhance Purposive Argumentation in the Value Judgment Formalism, ICAIL 2013 Proceedings [FORTHCOMING].
[12] Greenwood, K., Bench-Capon, T.J.M., McBurney, P., Towards a Computational Account of Persuasion in Law. ICAIL 2003 Proceedings, 22-31, ACM, 2003.
[13] G. Sartor. Doing justice to rights and values: teleological reasoning and proportionality. Artif. Intell. Law, 18(2):175¨215, 2010.

Legal argumentation with cases

Padmaja Sasidharan[1]

One of the objectives of the AI and Law community is to build legal expert systems. In this position paper I am going to take a stand on two main points - lack of legal expert systems that can be accessed by lay-users; focus on the application of state of the art AI techniques to overcome the hurdles in developing legal expert systems that can be accessed by lay-users e.g.natural language processing, common-sense reasoning.

The potential users for legal expert systems could be broadly classified in terms of depth of legal knowledge possessed by the user and their frequency of using the system: 1) people with legal knowledge who might want to use legal expert systems regularly to construct and evaluate their legal arguments e.g. lawyers, judges. 2) people who want to acquire legal knowledge and therefore use it regularly for a short period of time in their life e.g. students, people who work with legislative rules 3) people who work with legislative rules and therefore use legal expert systems in their day-to-day work life e.g. a person working with tax law system and 4) people who are probably going to use the system only a few times in their life to know about their legal situation in some dispute - lay-users. The existing legal expert systems are mostly applicable only to the first three groups of users.

Domains which have a well structured statute (e.g. British Nationality Act or Home office deduction from Federal Income tax law), reasoning starts at the statutory level; when there is a problem with the application or interpretation of a statutory norm to a fact, there is a transition from statutory reasoning to reasoning with precedents or purpose of law. The formalisation of statute and the problems associated with a legal expert system carrying out statutory reasoning are not new for the AI and Law community (e.g.[9] [4]). In domains where there are no well defined rules, the main means for assessing the legal situation is through factors. Factors are a collection of facts that have some legal significance in a given case. In these domains arguments are constructed and evaluated by comparing and distinguishing the current case and its precedents in terms of factors (e.g Trade secrets).

One of the major reasons for the limitations in the development of legal expert systems for lay-users is that legal expert systems for lay-users has always been under debate. Providing legal advise is a crucial task and involves factors such as interpretation of human actions, emotions, etc. These reasons will always be raised to question the practicality of legal expert systems for lay-users. But with a thorough research on the domain, potential users and identification of the barriers, a legal expert system designed for those users with the application of appropriate AI techniques to overcome the barriers, could be an essential one for that domain. The reason for the limited ground work on the development of legal expert systems that can be accessed by lay-users is that the research on representation of precedents and their interaction with statute and purpose of law have been restricted to a format that, can for the most part be worked on/with only by people who work regularly with legal materials. The representation of cases as factors and dimensions in a legal expert system were introduced in HYPO [1]. A factor is applied to a case depending on whether a certain pattern of facts are present in the case or not. Dimensions were used to indicate the extent to which a factor is present in the case. CATO[15] introduced a factor hierarchy using which two other argument moves can be made in addition to the argument moves in HYPO. Reasoning based on factors have been dominating in AI and Law since HYPO. Reasoning with cases has been described as a process of constructing, evaluating and applying a theory in [13]; theories explaining the decision in precedents can be created using the theory constructors; competing theories can be evaluated and the best theory can be applied to solve a new case. [14] gives an account of the CATE, AGATHA and ETHEL which aides in theory construction. The theory construction model is a powerful model but the application of such a system is more beneficial to people familiar with legal reasoning. In [2] a body of case-law were represented using Dung's abstract argumentation framework of [10]; the same body of case-law were then represented using value-based argumentation framework of [3] in [14] and extended argumentation framework of [11] in [12]. Representation of the case-law as an argumentation

[1]Supervisors:
Professor Andrew JI Jones, Kings College London.
Dr Jeroen Keppens, Kings College London.
Dr Claire Henderson, Kings College London.
Dr Graeme Lockwood, Kings College London.

framework enables the reconstruction of the reasoning that was carried out in precedents and also solve new cases using the constructed argumentation framework. Although each argument in these frameworks were abstract, most of them seem to take the form of factors and dimensions. In [14], BDI agents were used to generate the value-based argumentation frameworks and it was shown that the knowledge required by the agents to generate those argumentation frameworks can be provided by the theory construction tools. Legal expert systems built based on these models can be very useful for legal practitioners and people who wish to understand a particular legal domain. [7] provides an account of a model of how courts are constrained to respect precedents. The paper views precedential constraints as defeasible rules. This model again works with a set of factors and also mentions the lack of well understood work on assignment of factors to facts. In [6] illustrated the formalisation of CATO style arguments in ASPIC$^+$ framework. Again this paper indicates the lack of well accepted model for the assignment of factors to facts.

From my experience working on the EQUALS project, legal expert systems designated to deliver preliminary legal advice can be of great benefit in some domains such as the application of the *Equality Act 2010*. The aim of EQUALS project was to study the potential of legal decision aids in delivering employment related legal advice to mental health patients. To achieve this we formalised relevant sections from the Equality Act 2010, UK and developed a rule-based legal decision aid. The Equality Act 2010 concerns people with various protected characteristics such as age, race, disability, etc; it aims to prevent discrimination and promote equality. Our focus was specifically on "mental health problems" and "employment" [2]. We worked as a team with mental health experts[3]and an employment lawyer[4]. We formalised the relevant sections from the Equality Act 2010 to advise people on: whether their illness is covered by the Act (to be protected by the Act because of disability, the person's health condition should satisfy the conditions set out by the Act); whether they are being discriminated against at work because of their disability and what the legal consequences are; whether they are entitled to adjustments in the workplace; whether those adjustments are reasonable; what are the legal consequences if their employer fails to provide reasonable adjustments. The accuracy of the system was tested by comparing the conclusions made by the system for a set of cases against an employment lawyer's conclusions for those cases. We assessed and studied the usability and the desirability of the rule-based decision-aid based on feedback from potential users (mental health patients who had no knowledge about the Equality Act 2010) and professionals (vocational health advisers and occupational physicians who had some knowledge about the Equality Act 2010). The results of the user testing were not completely satisfactory owing to the dense language used in the legislation, reflected by the rules in the rule-based system. But the user feedback did not fail to indicate high desirability for such systems. The feedback from the potential professional users was very positive - again indicating that such systems are of more benefit to users with some knowledge about the legal domain. The professional users confirmed the desirability of such systems for lay-users by citing some factors. For instance the matters involved in these cases are very delicate and personal - they involve a person's mental health problem which are often not perceived as a health problem by the person; it involves details about their relationship with their employer and their colleagues. As a result some people may hesitate to discuss these matters with their vocational health advisers. So such users may prefer to use a software to get their advice which would assure them that their information is kept confidential.

As discussed earlier, the current state of the art in legal reasoning with cases mostly involves reasoning with precedents in terms of factors. Thus making the current AI and Law applications applicable mainly to legal practitioners or law students. This brings us to the widely discussed issue in case-based reasoning - "How to carry out facts-to-factor transformation"? Cases are represented in the form of factors. We need to focus on tools to bring about the facts-to-factors transformation. The EQUALS project was limited to testing a rule-based system; The next stage is my PhD in which I am working on a hybrid system that can support statutory interpretation. One such system was CABARET which achieved statutory and case-based reasoning [4]. While CABARET works on the basis of some control heuristics, I am working on a model that works on the basis of a classification of statutory interpretation problems; and when the type of interpretation

[2]The EQUALS project is now being maintained and expanded to include other protected characteristics by Monad Solutions (UK). The EQUALS project is funded by Guy's & St Thomas' Charity.

[3]Dr Claire Henderson and team, the Institute of Psychiatry, King's College London .

[4]Dr Graeme Lockwood, Department of Management, King's college London.

problem has been identified, case-based reasoning or teleological reasoning is applied, as appropriate. The basis for this classification is Prakken's description of three forms of open-texture - "Vagueness", "variable-standard" and "defeasibility"[5]. This has paved the way for the design of a rule-based system that can use a back-end support reasoner to perform statutory interpretation and therefore making a rule-based system accessible to lay-users. In [8] Ashley et al tried out assigning factors to case texts automatically using classifiers. In order to build a legal expert system for lay-users we need to go a step backward and start from collecting the fact situation. In my PhD thesis I am mainly focusing on ways to collect information from the user, transforming it into a coherent fact situation, and reasoning about the application of vague statutory predicates to that fact situation. As part of this I am also focusing on the representation of precedents and purpose of law in a form that can support reasoning with facts rather than factors. The proposed design and the addressing of these issues could be a positive step towards developing legal expert systems for lay-users.

References

[1] K.D. Ashley. *Modeling Legal Argument- reasoning with cases and hypotheticals*. Artificial Intelligence and legal reasoning. MIT Press, 1990.

[2] T.J.M. Bench-Capon. Representation of case-law as an argumentation framework. In *Legal Knowledge and Information Systems*, 2002.

[3] T.J.M. Bench-Capon and P.Dunne. Value based argumentation frameworks. In *Artificial Intelligence*, 2002.

[4] E.L.Rissland and D.B. Skalak. Cabaret: rule interpretation in a hybrid architecture. *Man-Machine Studies*, 34:839–887, 1991.

[5] H.Prakken. *Logical tools for modeling legal argument*. Kluwer Academic Publishers, 1993.

[6] T.J.M.Bench-Capon H.Prakken, A.Wyner and K.Atkinson. A formalisation of argumentation schemes for legal case-based reasoning in aspic+. In *Journal of Logic and Computation*, 2013.

[7] J.F.Horty and T.J.M.Bench-Capon. A factor-based definition of prcedential constraint. *Artificial Intelligence and Law*, 20:181–214, 2012.

[8] K.D.Ashley and S.Bruninghaus. Automatically classifying case texts and predicting outcomes. *Artificial Intelligence and Law*, 17:125–165, 2009.

[9] R.A.Kowalski-F.Kriwaczek P.Hammond M.J.Sergot, F.Sadri and H.T.Cory. The british nationality act as a logic program. *Communications of the ACM*, 29:370–386, 1986.

[10] P.M.Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-persons games. *Artificial Intelligence*, 77:321–357, 1995.

[11] S.Modgil. Reasoning about preferences in argumentation frameworks. *Artificial Intelligence*, 173:901–934, June 2009.

[12] 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.

[13] T.J.M.Bench-Capon and S.Modgil. Case-law in extended argumentation framework. In *ICAIL '09 Proceedings of the 12th International Conference on Artificial Intelligence and Law*, 2009.

[14] K.Atkinson T.J.M.Bench-Capon and A.Chorley. Persuasion and value in legal argument. *Journal of Logic and Computation*, 15:1075–1097, 2005.

[15] V.Aleven. Using background knowledge in case-based legal reasoning: A computational model and an intelligent learning environment. *Artificial Intelligence*, 150:183–237, 2003.

Robot, Esq.

Josh Blackman
Assistant Professor, South Texas College of Law

Advances in artificial intelligence are transforming many aspects of our society, from Google„s autonomous cars to IBM„s Watson defeating the Jeopardy! world champion. The legal profession, as well, is evolving from today„s time-consuming, customized labor-intensive legal market to tomorrow„s on-demand, commoditized law„s information revolution.

In the not-too-distant future, artificial intelligence systems will have the ability to reduce answering a legal question to the simplicity of performing a search. Imagine a program similar to the iPhone„s Siri app„ call it *Harlan*, your personalized virtual litigation assistant. A would-be litigator could tell *Harlan* about the case at hand: the relevant parties, the facts, the merits, the remedy sought, and share any relevant documents. Based on an advanced algorithm that mapped out the relationship between all of the relevant case law, statutes, and regulations, *Harlan* could generate forecasts of how the case would be resolved with different judges in different courts, and perhaps even recommend an ideal forum (call it fantasy-forum-shopping).

Harlan could explain how best to strategize the litigation, what types of motions would be most successful, and how to structure arguments. With advances in artificial intelligence, it is not difficult to conceive of *Harlan* even using document-assembly methods to draft the briefs (many sections of briefs today are copied from boilerplate), or at least check the persuasiveness of the arguments against other successful arguments already accepted by courts.

Harlan would also work wonders for non-lawyers. A person could download the app, talk to *Harlan* in plain-English, explain his or her problem, and listen to possible remedies. This process may or may not involve paying a lawyer. *Harlan* would improve access to justice.

As transformational as this technology may be, it raises fundamental questions about how we view our legal system, the representation of clients, and the development of our law. Before we proceed to develop, implement, and rely on this technology, we must first grapple with three important issues inherent in this change. First, what are the ethical implications of this technology to the traditional attorney-client relationship? Second, what are the jurisprudential implications of non-humans making and developing legal arguments? Third, how should we, or not, develop the legal and regulatory regimes to allow systems to engage in the practice of law?

Before considering whether we *can* develop *Harlan,* we must pause to consider whether we *should* develop *Harlan*? Will it actually improve conditions for attorneys, non-attorneys, and the rule of law? This article explores how advances in artificial intelligence will impact the practice of law, and lays out a framework that considers key issues with this important technology. This article begins the discussion of Robot, Esq.

I. Ethical Issues

Allowing *Harlan* to dispense legal advice without a human intermediary raises several very important questions. Would an attorney-client relationship be possible if a networked-distributed algorithm is used by many robots? What about the rules of confidentiality if the robot„s algorithms are improved by sharing and aggregating litigation strategies from other cases (think of how Google improves his search accuracy by discerning trends and patterns in usage)? What about conflicts of interest? If two opposing parties are both represented by *Harlan*, how would the algorithms handle that conflict? What about asking *Harlan* to do „the right thing„? Can we program the ethos of Atticus Finch? How would these systems embody zealous advocacy and representation? Would Harlan have an obligation to report unethical conduct by a client? Would *Harlan* withdraw under the circumstances where a real lawyer would withdraw? How would this technology be used to promote access to justice, and provide representation to indigent clients?

Today, predictive coding algorithms are already replacing document review attorneys. If clients become accustomed to *Harlan* providing instant, customized answers, the desire to procure attorneys may be further diminished. This shift in demand will result in changing the structure of the legal profession, and modifying the workforce. How will people react to robots taking jobs once reserved for humans?

II. Jurisprudential Issues

Beyond the ethical considerations, attorneys must confront what it would mean to have computer systems arguing, and perhaps even resolving cases or controversies. A primary concern is the potential for bias. Algorithms are not transparent. How Google orders search results narrowly avoided an antitrust suit by the FTC. The ability of these algorithms, perhaps influenced by biases,, overt or implicit,, will have a great impact on what we see and think. Applied to the law, the risk of bias in an algorithm could be pernicious. It would be quite easy,, and lucrative,, for certain interests to capture the algorithm and make the results skew in one direction. The very transparency that is the sine qua non of the adversarial process would have to apply to this technology in order to give it any legitimacy of unbiased assisted-decision-making.

Another possible problem is the potential ossification of the law. If a system is simply producing the best argument based on previous precedents (especially if that was a winning argument), the precedents will not evolve and change. Courts, being fed the same arguments over and over again, will have less space to advance the jurisprudence. This iterative process can result in a legal stagnation. Courts that already reuse boilerplate language in unpublished orders are already contributing to this ossification. We would need to consider how this technology impacts our fundamental notions of fairness and due process,, and how courts would respond to this formulaic recitation of the same arguments over and over again. Many flesh-and-blood jurists may reject these positions to assert judicial independence from predictive algorithms.

III. Legal and Regulatory Issues

The final issue is likely to be the first problem confronted,, *can* computers solve legal problems. Although the ethical and jurisprudential implications are significantly more important, developers and technologists are already forging ahead with this technology, and are on a collision course with a number of legal and regulatory regimes that will serve as barriers to the proliferation of this technology.

First and foremost, this technology will have to grapple with state unauthorized practice of law (UPL) regimes. Today in the United States, the practice of law is regulated by state bar associations. The definition of engaging in the practice of law is quite vague, and ill-defined. While early iterations of this technology are unlikely to be challenged, future, more sophisticated algorithms that can dispense legal advice may constitute practicing law. Bar associations and attorneys will challenge such programs as engaging in the unauthorized practice of law and try to shut them down,, similar to the suits against LegalZoom in the United States.

This regulatory issue is not limited to the practice of law. Nurses, doctors, architects, professional engineers, and a host of other regulated professions,, all subject to various occupational licensing regimes, and all professions that can be automated,, will need to contend with the specter of robots performing these tasks. Entrenched interests will, to some degree, avail themselves of the regulatory arm of the state to block robotic competition. These dynamics will apply in the United States, and around the world.

Second, issues of liability are quite uncertain. What happens if *Harlan* gives bad legal advice? Would a product liability suit, or malpractice suit lie? If so, against whom? The developer of the software? Would *Harlan* obtain malpractice insurance? Who would insure that? Would *Harlan* be subject to malpractice in the same fashion an attorney would? What if *Harlan* prepares an invalid document that results in material losses? Liability analyses for autonomous cars provides relevant frameworks to consider these issues.

This article opens the first chapter in this process of building Robot, Esq., and sets forth an agenda of issues to consider as the intersection between law, technology, and justice merges.

Reasoning with Cases

Selected Bibliography

Trevor Bench-Capon
Department of Computer Science
University of Liverpool

General History of the Field

Trevor J. M. Bench-Capon, Michal Araszkiewicz, Kevin D. Ashley, Katie Atkinson, Floris Bex, Filipe Borges, Danièle Bourcier, Paul Bourgine, Jack G. Conrad, Enrico Francesconi, Thomas F. Gordon, Guido Governatori, Jochen L. Leidner, David D. Lewis, Ronald Prescott Loui, L. Thorne McCarty, Henry Prakken, Frank Schilder, Erich Schweighofer, Paul Thompson, Alex Tyrrell, Bart Verheij, Douglas N. Walton, Adam Zachary Wyner: A history of AI and Law in 50 papers: 25 years of the international conference on AI and Law. Artif. Intell. Law 20(3): 215-319 (2012)

Thorne McCarty

L. Thorne McCarty, Reflections on TAXMAN: An Experiment in Artificial Intelligence and Legal Reasoning, *Harvard Law Review*, vol. 90, pages 837-893 (1977).

L. Thorne McCarty, N. S. Sridharan: The Representation of an Evolving System of Legal Concepts: II. Prototypes and Deformations. IJCAI 1981: 246-253

"A Computational Theory of Eisner v. Macomber," in C. Ciampi, ed., *Artificial Intelligence and Legal Information Systems*, pages 329-355 (North-Holland Publishing Co., 1982).

L. Thorne McCarty: An Implementation of Eisner v. Macomber. ICAIL 1995: 276-286

Edwina Rissland

Edwina L. Rissland: Examples in Legal Reasoning: Legal Hypotheticals. IJCAI 1983: 90-93

Edwina L. Rissland, Jody J. Daniels: A Hybrid CBR-IR Approach to Legal Information Retrieval. ICAIL 1995: 52-61

Edwina L. Rissland, M. Timur Friedman: Detecting Change in Legal Concepts. ICAIL 1995: 127-136

HYPO

Edwina L. Rissland, Kevin D. Ashley: A Case-Based System for Trade Secrets Law. ICAIL 1987: 60-66

Kevin D. Ashley, Edwina L. Rissland: A Case-Based Approach to Modeling Legal Expertise. IEEE Expert 3(3): 70-77 (1988)

Kevin D. Ashley: Reasoning with Cases and Hypotheticals in HYPO. International Journal of Man-Machine Studies 34(6): 753-796 (1991)

Ashley, K.D., (1990). *Modeling Legal Argument: Reasoning with Cases and Hypotheticals*. The MIT Press / Bradford Books, Cambridge, MA.

CABARET

Edwina L. Rissland, David B. Skalak: CABARET: Rule Interpretation in a Hybrid Architecture. International Journal of Man-Machine Studies 34(6): 839-887 (1991)

Edwina L. Rissland, David B. Skalak: CABARET: Rule Interpretation in a Hybrid Architecture. International Journal of Man-Machine Studies 34(6): 839-887 (1991)

David B. Skalak, Edwina L. Rissland: Arguments and cases: An inevitable intertwining. Artif. Intell. Law 1(1): 3-44 (1992)

BANKXX

Edwina L. Rissland, David B. Skalak, M. Timur Friedman: BankXX: Supporting Legal Arguments through Heuristic Retrieval. Artif. Intell. Law 4(1): 1-71 (1996)

Edwina L. Rissland, David B. Skalak, M. Timur Friedman: Evaluating a Legal Argument Program: The BankXX Experiments. Artif. Intell. Law 5(1-2): 1-74 (1997)

CATO

Vincent Aleven, Kevin D. Ashley: Doing Things with Factors. ICAIL 1995: 31-41

Vincent Aleven, Kevin D. Ashley: Evaluating a Learning Environment for Case-Based Argumentation Skills. ICAIL 1997: 170-179

Aleven, V. (1997). Teaching Case-Based Argumentation Through a Model and Examples. Ph.D. Thesis, Intelligent Systems Program, University of Pittsburgh.

IBP

Stefanie Brüninghaus, Kevin D. Ashley: Predicting Outcomes of Case-Based Legal Arguments. ICAIL 2003: 233-242

A Predictive Role for Intermediate Legal Concepts, Ashley, K. and Brüninghaus, S. (2003). In the *Proceedings 16th Annual Conference on Legal Knowledge and Information Systems,* Jurix-03. pp. 153-162. Utrecht, The Netherlands. December. IOS Press, Amsterdam.

Stefanie Brüninghaus, Kevin D. Ashley: Generating Legal Arguments and Predictions from Case Texts. ICAIL 2005: 65-74

Stefanie Brüninghaus, Kevin D. Ashley: Generating Legal Arguments and Predictions from Case Texts. ICAIL 2005: 65-74

Reconstructing Factor Based Reasoning

T.J.M. Bench-Capon, *Arguing with Cases,* in Proceedings of JURIX 97, GNI Nijmegen, 1997, pp 85-100.

Henry Prakken, Giovanni Sartor: Modelling Reasoning with Precedents in a Formal Dialogue Game. Artif. Intell. Law 6(2-4): 231-287 (1998)

M. Allen, Trevor J. M. Bench-Capon, Geof Staniford: A Multi-Agent Legal Argument Generator. DEXA Workshop 2000: 1080-1086

John F. Horty, Trevor J. M. Bench-Capon: A factor-based definition of precedential constraint. Artif. Intell. Law 20(2): 181-214 (2012)

Trevor J. M. Bench-Capon: Representing Popov v Hayashi with dimensions and factors. Artif. Intell. Law 20(1): 15-35 (2012)

Argumentation

T.J.M. Bench-Capon, 2002. Representation of Case Law as an Argumentation Framework. In T. Bench-Capon, Daskalopoulu, A., and Winkels, R., (eds) Legal Knowledge and Information Systems, Proceedings of Jurix 2002. IOS Press: Amsterdam. pp 103-112.

Adam Zachary Wyner, Trevor J. M. Bench-Capon: Modelling Judicial Context in Argumentation Frameworks. J. Log. Comput. 19(6): 941-968 (2009)

Colen, S., Cnossen, F., & Verheij, B. (2009). How Much Logical Structure is Helpful in Content-Based Argumentation Software for Legal Case Solving?

Henry Prakken: Reconstructing Popov v. Hayashi in a framework for argumentation with structured arguments and Dungean semantics. Artif. Intell. Law 20(1): 57-82 (2012)

Values

Trevor J. M. Bench-Capon: The missing link revisited: The role of teleology in representing legal argument. Artif. Intell. Law 10(1-3): 79-94 (2002)

Henry Prakken: An exercise in formalising teleological case-based reasoning. Artif. Intell. Law 10(1-3): 113-133 (2002)

Giovanni Sartor Teleological arguments and theory-based dialectics. . Artif. Intell. Law 10(1-3): 95-112 (2002)

Trevor J. M. Bench-Capon, Giovanni Sartor: A model of legal reasoning with cases incorporating theories and values. Artif. Intell. 150(1-2): 97-143 (2003)

Trevor J. M. Bench-Capon, Katie Atkinson, Alison Chorley: Persuasion and Value in Legal Argument. J. Log. Comput. 15(6): 1075-1097 (2005)

Adam Zachary Wyner, Trevor J. M. Bench-Capon, Katie Atkinson: Arguments, Values and Baseballs: Representation of Popov v. Hayashi. JURIX 2007: 151-160

Trevor J. M. Bench-Capon, Sanjay Modgil: Case law in extended argumentation frameworks. ICAIL 2009: 118-127

Argumentation Schemes

Katie Greenwood, Trevor J. M. Bench-Capon, Peter McBurney: Towards a Computational Account of Persuasion in Law. ICAIL 2003: 22-31

Katie Atkinson, Trevor J. M. Bench-Capon, Peter McBurney: Arguing about cases as practical reasoning. ICAIL 2005: 35-44

Adam Zachary Wyner, Trevor J. M. Bench-Capon: Argument Schemes for Legal Case-based Reasoning. JURIX 2007: 139-149

H. Prakken, A.Z. Wyner, T.J.M. Bench-Capon and K.M. Atkinson (2013), A formalisation of argumentation schemes for legal case-based reasoning in ASPIC+. *Journal of Logic and Computation*, in press.

T.J.M. Bench-Capon, H. Prakken, A.Z. Wyner and K.M. Atkinson (2013). Argument Schemes for Reasoning with Legal Cases Using Values. Proceedings of *Fourteenth International Conference on Artificial Intelligence and Law* (ICAIL 2013). Forthcoming.

AGATHA

Alison Chorley, Trevor J. M. Bench-Capon: AGATHA: Using heuristic search to automate the construction of case law theories. Artif. Intell. Law 13(1): 9-51 (2005)

Theory Construction

Alison Chorley, Trevor J. M. Bench-Capon: An empirical investigation of reasoning with legal cases through theory construction and application. Artif. Intell. Law 13(3-4): 323-371 (2005)

Adam Zachary Wyner, Trevor J. M. Bench-Capon, Katie Atkinson: Towards formalising argumentation about legal cases. ICAIL 2011: 1-10

Neural Networks

Trevor J. M. Bench-Capon: Neural Networks and Open Texture. ICAIL 1993: 292-297

Semantic Networks

Karl Branting: Representing and Reusing Explanations of Legal Precedents. ICAIL 1989: 103-110

Karl Branting: Building Explanations from Rules and Structured Cases. International Journal of Man-Machine Studies 34(6): 797-837 (1991)

Karl Branting: A Reduction-Graph Model of Ratio Decidendi. ICAIL 1993: 40-49

Karl Branting: A computational model of ratio decidendi. Artif. Intell. Law 2(1): 1-31 (1993)

Extending Factors

Roth, B., & Verheij, B. (2004). Dialectical Arguments and Case Comparison. *Legal Knowledge and Information Systems. JURIX 2004: The Seventeenth Annual Conference* (ed. Gordon, T.F.), 99-108. Amsterdam: IOS Press.

Evidence

F.J. Bex, H. Prakken & B. Verheij (2007) Formalising argumentative story-based analysis of evidence. *Proceedings of the 11th International Conference on Artificial Intelligence and Law*, 1-10. New York: ACM Press.

F.J. Bex, S.W. van den Braak, H. van Oostendorp, H. Prakken, H.B. Verheij & G.A.W. Vreeswijk (2007) Sense-making software for crime investigation: how to combine stories and arguments? *Law, Probability & Risk* 6:145-168

F.J. Bex, P.J. van Koppen, H. Prakken & B. Verheij (2010) A Hybrid Formal Theory of Arguments, Stories and Criminal Evidence. *Artificial Intelligence and Law* 18:2, 123-152

F.J. Bex (2011) *Arguments, Stories and Criminal Evidence: A Formal Hybrid Theory*. Springer, Dordrecht.