A Default-Logic Framework for Legal Reasoning in Multiagent Systems

Vern R. Walker
Professor of Law
Hofstra University School of Law
121 Hofstra University
Hempstead, NY 11549
Vern.R.Walker@Hofstra.edu

Abstract

Using law and evidence to achieve fair and accurate decisions in numerous legal cases requires a complex multiagent system. This paper discusses a framework based on many-valued, predicate, default logic that successfully captures legal knowledge, integrates and evaluates expert and non-expert evidence, coordinates agents working on different legal problems, and evolves the knowledge model over time. The graphical syntax and the semantics of this framework allow the automation of key tasks, and the emergence of dynamic structures for integrating human and non-human agents. The logical basis of the framework ensures its applicability to knowledge and problem domains of similar complexity to law.

Introduction

This paper discusses a default-logic framework for capturing, integrating, applying, evolving, and coordinating the use of any practical knowledge in multiagent systems. Practical knowledge is default reasoning that uses inference structures, together with the available evidence, to warrant presumptive but defeasible conclusions about possible actions (Brewka, Dix, and Konolige 1997; Josephson and Tanner 1996; Kyburg and Teng 2001; Levi 1996; Pollock 1990; Prakken 1997; Toulmin, Rieke, and Janik 1984; Walton 1996, 2002). The knowledge domain discussed in this paper is legal knowledge, although the framework is applicable to decisional knowledge in any domain (Russell and Norvig 2003; Singh, Rao, and Georgeff 1999). The major engineering problems in this and similar domains are:

capturing knowledge in a structure that is suited for its content and use;

integrating expert and non-expert knowledge into a single practical model;

applying the knowledge model to solve the next problem;

evolving and adapting the knowledge model over time, on the basis of experience with solving problems; and

coordinating the use of such knowledge among autonomous agents (both human and non-human) and among different problem sets over time.

This paper outlines a particular default-logic framework (D-L Framework) for solving these problems. The D-L Framework utilizes a many-valued, default predicate logic adapted to model both rule-based and case-based reasoning in law. The Decision Apprentice™ software (DA software) of Apprentice Systems, Inc. incorporates the D-L Framework and has successfully applied it to law in the Legal Apprentice™ software.

Syntax, Ontology, and Semantics of the D-L Framework

The knowledge-capture environment of the D-L Framework can take various syntactic forms, but a graphical syntax appears easiest for human agents to use. Microsoft Office Visio™ shapes represent the elements of the ontology, and the shapes are defined to allow only permissible combinations of elements. Figure 1 shows representative shapes for the ontology discussed below. The DA software builds the knowledge model as the Visio shapes are selected, dragged, and connected.

The ontology and semantics for the D-L Framework is as follows:

Definite subjects: specific individuals named by proper names or definite descriptions (e.g.: Vern Walker; the paper being submitted to AAAI) (Chierchia and McConnell-Ginet 2000; Larson and Segal 1995; Rodes and Pospesel 1997);

Indefinite subjects: groups of one or more individuals identified solely by their attributes (e.g.: Americans over age 50; submissions to AAAI symposia) (Chierchia and McConnell-Ginet 2000; Larson and Segal 1995; Rodes and Pospesel 1997);

Predicates: propositional functions that generate propositions when supplied with the appropriate number of definite or indefinite subjects (e.g.: "… is a citizen of the United States"; "… is the author of … ") (Chierchia and McConnell-Ginet 2000; Copi and Cohen 1998; Larson and Segal 1995; Rodes and Pospesel 1997; Saeed 2003; Sainsbury 1991);

Concepts: categories used to classify indefinite subjects and predicates (e.g.: human being; ownership relations);

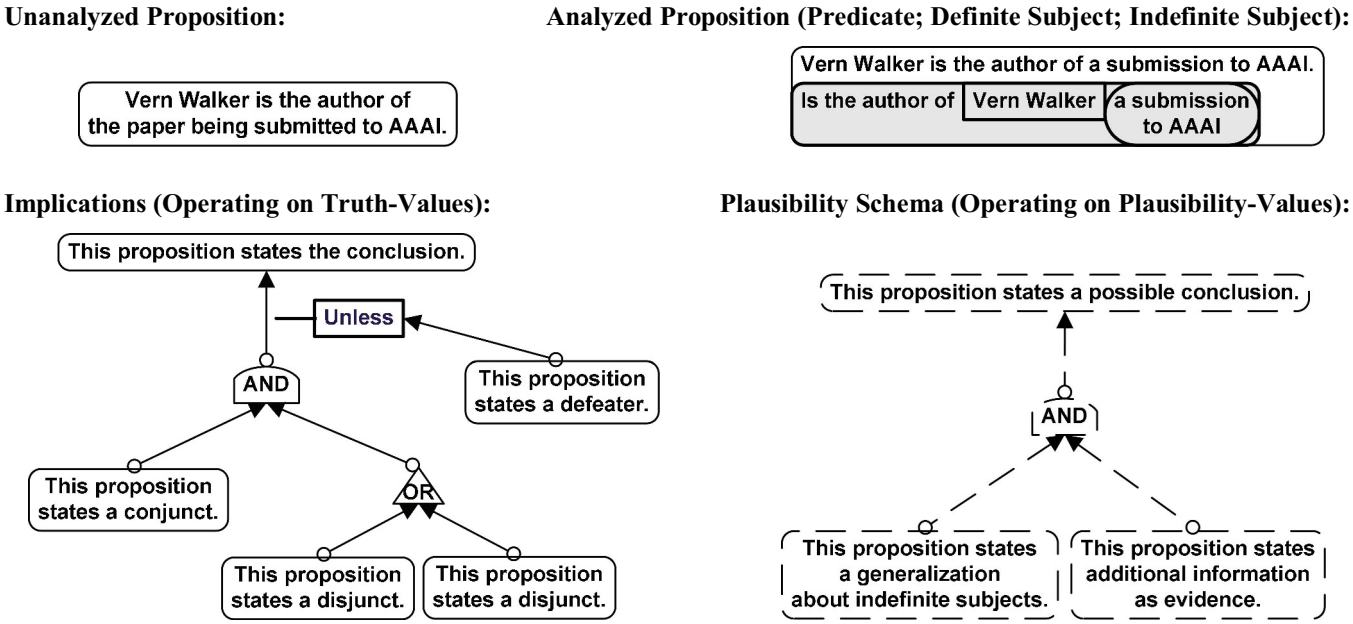

Figure 1. Illustrative Shapes for Selected Elements in the D-L Framework

Propositions: the informational content of declarative sentences or assertions, capable of having either a truth-value or a plausibility-value (e.g.: "Vern Walker is a citizen of the United States"; "Vern Walker is the author of the paper being submitted to AAAI"); a proposition can be either unanalyzed or analyzed into its predicate-subject structure, to the extent needed to warrant inferences (see Table 1 for illustration);

Many-valued, truth-functional connectives: the D-L Framework uses three (see Table 1 for illustration):

- Conjunction ("AND"): a connective for two or more propositions, and whose truth-value or plausibility-value equals the lowest such value among the conjunct propositions (Copi and Cohen 1998; Gottwald 2001; Rodes and Pospesel 1997; Sainsbury 1991);
- Disjunction ("OR"): a connective for two or more propositions, and whose truth-value or plausibility-value equals the highest such value among the disjunct propositions (Copi and Cohen 1998; Gottwald 2001; Rodes and Pospesel 1997; Sainsbury 1991);
- Defeater ("UNLESS"): a connective that sets the truth-value or plausibility-value of the conclusion equal to the inverse of the value of the defeater proposition whenever the defeater proposition has a positive truth-value or plausibility-value; otherwise, the truth-value or plausibility-value of the conclusion remains what it would have been in the absence of a defeater proposition (Brewka, Dix, and Konolige 1997; Pollock 1990);

Implications: conditional propositions consisting of a condition (composed of one or more connected propositions) and a conclusion, in which the *truth-value* of the conclusion is determined by the connective operating on the truth-values of the propositions in the condition; in the D-L Framework, the conclusion is graphed at the top, supported by its condition (see Table 1 for illustration); for an explanation of truth-values, see the paragraph below;

Entailments: implications that are "local" in the sense that the condition consists of a small number of completely identified propositions; entailments state the necessary and/or sufficient conditions for using concepts and predicates;

Plausible inferences: conditional propositions consisting of a condition (composed of one or more connected propositions) and a conclusion, in which the *plausibility-value* of the conclusion is determined by the connective operating on the plausibility-values in the condition (see Table 1 for illustration); for an explanation of plausibility-values, see the paragraph below;

Plausibility schemas: inverted directed acyclic schemas that produce plausible inferences in a particular case whenever evidentiary assertions having the appropriate logical form and a positive plausibility-value are substituted into the schemas;

Inference trees: inverted directed acyclic graphs consisting of chained levels of implications and plausible inferences, with (1) the ultimate conclusion at the top, (2) the upper branches consisting of implications (the *implication tree* or rule-based region of the

inference tree), and (3) lower levels of branches consisting of plausible inferences (the plausibility-schema region of the inference tree).

Truth-values are attributes of propositions that take one of three values: "true / undecided / false." When reasoning begins in a particular situation, the truth-values of the conclusions and conditions within the applicable implication tree are all undecided. Attaching and evaluating evidentiary propositions may then change the truth-values of particular conditions, which may in turn change the truth-values of conclusions.

Plausibility-values are attributes of propositions that take a value from a plausibility scale. Plausibility scales can have any number of values, either qualitative or quantitative. For example, a qualitative plausibility scale might be ordinal and have five values (such as "true / probably true / undecided / probably false / false") or seven values (such as "almost certainly true / highly likely / probably true / undecided / probably false / highly unlikely / almost certainly false"). By contrast, conventional mathematical probability is an infinite-valued quantitative plausibility scale, using the set of real numbers between zero and one and having values such as 0.56. When evaluating evidentiary propositions in a particular case, an agent selects a suitable plausibility scale for each particular proposition and assigns a plausibility-value from that scale to the evidentiary proposition.

Implication Trees As Capturing Rule-Based Knowledge

The D-L Framework is designed to model, for any particular legal case, the reasoning that warrants the legal findings, decisions, and actions in that case. The D-L model for that reasoning is an inference tree.

The upper portion of any inference tree is an implication tree, which models all of the acceptable implications or lines of reasoning. The ultimate conclusion at the top roots a tree structure because lower-level conditions never depend for their truth-values on a higher-level proposition in the same branch. Implication trees branch downward and outward from a single root conclusion. For example, the rules of tort law for battery, which justify a court judgment that the defendant must pay damages, can be modeled as one large implication tree that begins as shown in Figure 2. The legal interpretation of this tree is that "the defendant is liable to the plaintiff for battery" (conclusion) if (1) "the defendant performed a voluntary act," (2) "the defendant acted intending to cause a harmful or offensive contact with a person," and (3) "the defendant's act caused a harmful or offensive contact with the plaintiff," but this line of reasoning is defeated if "the defendant was privileged to perform the action," which would be true if either "the defendant acted reasonably in making a lawful arrest" or "the defendant acted reasonably in defending herself from intentionally inflicted bodily harm" (American Law Institute 1965; Dobbs 2000). The bottom propositions of each branch of an implication tree, where the legal rules end, are the *terminal propositions*. The truth-value of a terminal proposition can be determined only by stipulation or by the schematized evidentiary assertions that are attached in a particular case and then evaluated.

In addition to implication trees, the D-L Framework uses entailments to model local semantic rules. Any proposition in an implication tree can be analyzed into a predicate and one or more subjects. Entailments are "local" because the truth-value of the conclusion is determined by the truth-values of a specifiable set of conditions. Some entailments identify class/subclass relationships among concepts, so that any attribute of class members is a *necessary* attribute of subclass

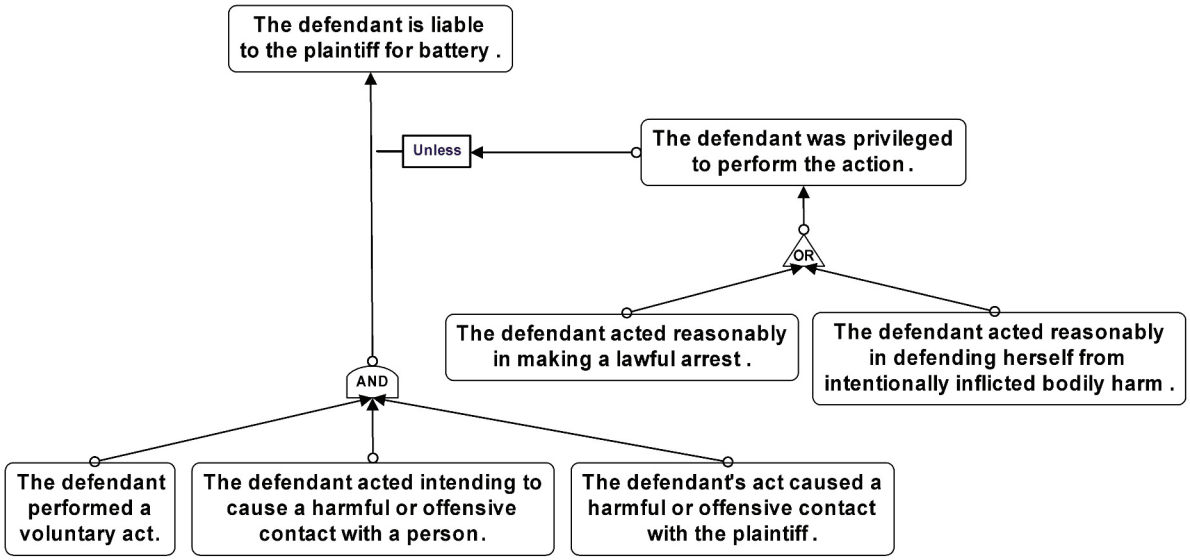

Figure 2. Implication Tree for the Tort Rules of Battery

members. Other entailments state a set of jointly *sufficient* conditions for a subject's satisfying a particular predicate. A *definition* combines both types of relationship (a single set of necessary and jointly sufficient conditions) into a single statement of equivalence. Entailments are useful in warranting inferences between propositions in different branches of the same implication tree, or between propositions within different implication trees. A dictionary of entailments can also operate across many knowledge domains. Such local rules, however, only supplement the strategic work of an implication tree, which is designed to model all of the rules that are relevant to proving a particular ultimate conclusion.

Plausibility Schemas as Applying and Integrating Knowledge

Reasoned decision-making in a particular situation involves attaching schematized evidence to one or more of the terminal propositions of an implication tree, evaluating the plausibility-values of the propositions in that evidence (the *evidentiary propositions*), and using those plausibility-values to assign truth-values to terminal propositions, which logical connectives can then use to propagate truth-values up the implication tree.

Choosing the plausibility scale to employ for evaluating any particular evidentiary assertion depends upon the pragmatic context – that is, upon the precision needed in the content and upon the potential for error acceptable in assessing plausibility (Walker 2004). For example, some tasks require only measurements in inches and accept even a moderate degree of plausibility for decision-making, with the result that even a single measurement with an ordinary ruler will yield sufficiently accurate values. Some of NASA's tasks, however, may require measurements in microns and a high level of quantitative plausibility. In general, as the level of required precision increases, the potential for error inherent in assessing plausibility also increases, as well as the cost of producing sufficient evidence to make the ultimate conclusion sufficiently plausible. Legal decision-makers try to use plausibility scales that achieve an acceptable balance in the pragmatic context.

Optimal decision-making in law, as in many practical fields, often requires the evaluation of both expert and non-expert evidence, using both quantitative and qualitative scales of plausibility. Plausibility schemas are patterns of evidentiary propositions that use the same three connectives as implication trees, but with the connectives operating on the plausibility-values of the evidentiary propositions to assign a plausibility-value to the conclusion. Figure 3 illustrates a D-L Framework logic diagram for schematizing the evidence supporting a legal finding. The dashed lines in the shapes indicate that the evaluation is operating on plausibility-values, not on truth-values. The interpretation of such a schema is that if the evidentiary propositions on the lower levels are plausible, then the conclusion at the top is plausible (or implausible) as determined by the plausibility connectives. Plausibility schemas therefore model default reasoning to conclusions that are presumptively plausible or implausible, but still defeasible.

Since agents may adopt different plausibility scales for evaluating different evidentiary propositions, there must be a rule for operating on a mixture of plausibility scales – for example, where one conjunct of a condition has a plausibility-value on a seven-point ordinal scale and another conjunct in the same condition has a quantitative value on the real-number scale. For conjunction and disjunction, such a rule requires determining whether a particular value on one scale is lower (for conjunction) or higher (for disjunction) than a value on another scale. After such an ordering of values, the schema can evaluate the conclusion on the plausibility scale of the critical evidentiary assertion – that is, for conjunction, the evidentiary proposition with the lowest plausibility-value, and for disjunction, the evidentiary proposition with the highest plausibility-value.

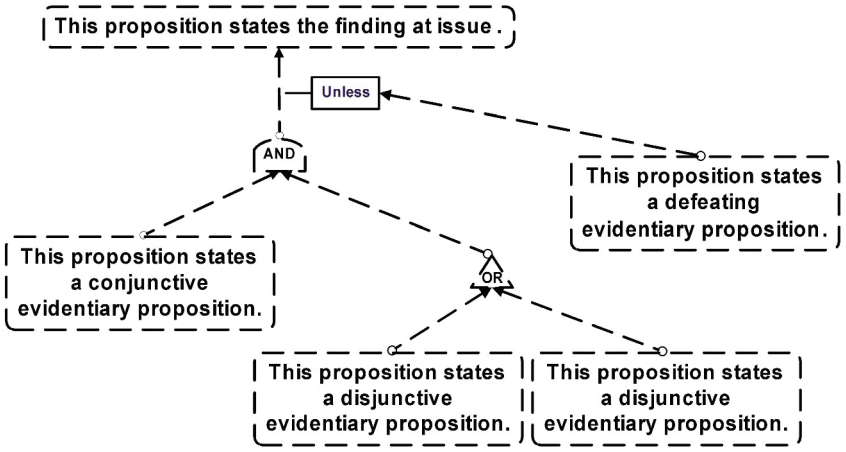

Figure 3. Illustration of a Plausibility Schema in the D-L Framework

In the case of a plausibility defeater, if the defeater proposition has a positive plausibility-value, then the defeater connective assigns to the conclusion the degree of plausibility that is the inverse to the plausibility-value of the defeater proposition. That is, as the plausibility of the defeater proposition increases, the plausibility of the conclusion decreases (alternatively, the *im*plausibility of the conclusion increases). For example, on the seven-point plausibility scale above, if the plausibility-value of the defeater proposition is "highly likely," then the plausibility-value of the conclusion would be "highly unlikely"; on a plausibility scale of mathematical probability, if the defeater's plausibility-value is 0.56, then the conclusion's plausibility-value would be 0.44 (1 – 0.56).

One component of perhaps all plausibility schemas is a *generalization* (Chierchia and McConnell-Ginet 2000; Copi and Cohen 1998; Kadane and Schum 1996; Rodes and Pospesel 1997; Schum 1994; Toulmin 1958). A generalization is a proposition that usually asserts that it is true in some situations but not all situations. When a generalization is analyzed into a predicate and one or more indefinite subjects, then it usually asserts that its predicate accurately describes only a proper subset of an indefinite subject class. Examples of generalizations are: "most witnesses testifying under oath tell the truth"; "one-third of Americans are overweight"; and "60% of the test group in the study developed the disease." These generalizations have the following logical forms (respectively): "most As are Bs"; "X/Y of As are Bs"; and "X% of the members of group A are members of group B." Logicians call group A the "reference class" or "reference group" for the generalization (Kyburg 1990). Two logical attributes of a generalization that can affect its plausibility-value are its degree of quantification and any modal hedge employed. Generalizations imply or explicitly assert a degree of quantification over the reference class – that is, they characterize the portion of A that is asserted to be B. Moreover, generalizations often contain an explicit modal "hedge" that qualifies the entire assertion. Examples of modal hedges are expressions of frequency (e.g., "sometimes" or "often"), typicality (e.g., "typically" or "normally"), temporal limitation (e.g., "in the past" or "at least for the immediate future"), or degree of confidence of the speaker (e.g., "perhaps" or "almost certainly"). Generalizations may derive from scientific, technical or other specialized knowledge, or they derive from personal experience or "common sense." Therefore, generalizations in plausibility schemas may represent either expert or non-expert conclusions of fact.

The two primary factors in selecting which plausibility schema to use in reasoning to a particular terminal proposition are (1) the logical form of the terminal proposition and (2) the nature of the available evidence. First, whether the terminal proposition is a generalization about a group (indefinite subject) or a proposition about a specific individual (definite subject) will determine what kind of schema is allowed. The D-L Framework only allows the use of schemas whose conclusions have a logical form identical to that of the terminal proposition. Second, evidence that is scientific and statistical would be schematized differently than eyewitness testimony. The agent making the decision selects a schema that fits the terminal proposition and the evidence in the particular case. This means that the schematized evidence is specific to the particular case, whereas the implication tree is generic to all cases within the knowledge domain.

Finally, in order for a plausibility schema to provide an inference from plausible evidence to a decision, there must be a rule for determining the truth-value of a terminal proposition as a function of the plausibility-value of the schematized evidence attached to that proposition. In legal terminology, this rule is the applicable "standard of proof" (James, Hazard, and Leubsdorf 1992; Walker 1996). For example, the standard of proof for most issues of fact in civil cases is preponderance of the evidence. Under this rule, if the schema evaluates the totality of attached evidence as having any plausibility-value other than "undecided," then the schema assigns the corresponding truth-value to the terminal proposition – that is, it assigns the value "true" to the terminal proposition if the schema evaluates the attached evidence as plausible to any degree, and assigns the value "false" to the terminal proposition if it evaluates the attached evidence as implausible to any degree. Standards of proof are the links between the output of schematized, evaluated evidence and the input to an implication tree.

An example of a particular plausibility schema is the *direct-inference schema*, which models one type of reasoning that warrants a conclusion about a definite subject (Kyburg 1983; Levi 1977, 1981; Pollock 1990; Salmon 1973). Examples of such conclusions are "the tire that caused the accident had a defect" and "the claimant contracted pneumoconiosis," where "the tire" and "the claimant" are definite subjects (Director, Office of Workers' Compensation Programs, Department of Labor v. Greenwich Collieries; Kumho Tire Co. Ltd. v. Carmichael). A D-L Framework diagram for the plausibility schema for direct inference is shown in Figure 4. The plausibility connective "AND" conjoining the branches assigns a plausibility-value to the conclusion that is equal to the plausibility-value of the least plausible of these four conjuncts.

In the plausibility schema shown in Figure 4, the four conjuncts state the evidentiary propositions that render the conclusion plausible. The first evidentiary proposition (from the left) is a generalization asserting that most members of category A are also members of category B. The second evidentiary proposition asserts that the specific individual that is the definite subject of the conclusion (S) is a member of category A. The third evidentiary proposition asserts that the definite subject S is a random member of A with respect to being B. This reflects the reasoning that if S is drawn from A in a simple random manner, then the probability that S is a B will approximate the ratio of the number of Bs in A to the total number of As. Finally, as a fourth factor in the reasoning, the degree of quantification in the generalization limits the range of probabilities that can plausibly be asserted in a conclusion. In this schema, the quantification in the generalization that "most" As are Bs fits the assertion in the

conclusion that "probably" S is a B. However, if only 10% of As were Bs, then this evidence could not warrant a plausible conclusion that S is "probably" a B.

The direct-inference plausibility schema also suggests how plausibility schemas can integrate expert evidence with non-expert evidence, and quantitative information with qualitative information. In a direct-inference schema adapted for statistical evidence, the generalization might state a percentage degree of quantification (e.g., 60%) and have a high level of plausibility. At the same time, there may be substantial uncertainty about whether the specific individual S is in fact a member of A. In such a case, the low plausibility-value for the second evidentiary proposition might be the critical minimum value for the conjunction, resulting in a correspondingly low plausibility-value for the conclusion. The evaluating agent might then have several strategies available for increasing the plausibility-value of that second evidentiary proposition. Alternatively, the agent might rely on a different line of reasoning altogether, using different schematized evidence, thus bypassing this weak direct-inference evidence.

The branches of a plausibility schema can themselves generate a chain of plausibility schemas, with the evidentiary propositions of one schema becoming the conclusions of lower-level plausibility schemas. For example, the second evidentiary proposition of the direct-evidence schema in Figure 4 is itself a proposition about a definite subject, so in a particular case such an assertion could become the conclusion of additional direct-inference evidence. At some point in any particular branch, however, an evaluating agent must stipulate plausibility-values for evidentiary propositions – using intuition (human agents), default values (human and non-human agents), sensitivity analysis (all agents), or some other method.

Process Rules for Evolving Knowledge Models and Coordinating Their Application

Process rules govern the dynamics of default reasoning within multiagent systems. They coordinate the process of applying implication trees and plausibility schemas in a particular case, and also the process of evolving new trees and schemas. Process rules therefore play an important role in emergent reasoning and behavior.

Legal systems use process rules to evolve new legal rules from past legal decisions, and to evolve new plausibility schemas from past evaluations of evidence. Legal systems also use such rules to coordinate multiple courts and factfinders, in an effort to achieve reasonably accurate and consistent decision-making.

A key feature of the D-L Framework is that it requires no new types of logical structures to model legal process rules. This means that the D-L Framework can capture the domain knowledge for evolution and coordination, and integrate that knowledge into the same model that captures substantive legal rules and the factfinding in particular cases. While the modeling of legal process rules provides an area for more research, the D-L Framework offers a very promising approach, for the reasons discussed in this section of the paper.

Evolving Implication Trees

The legal reasoning behind a decision to adopt new legal rules (i.e., to modify implication trees) balances "policy rationales" for and against a rule change. Policy rationales can be either epistemic or non-epistemic (Walker 2003, 2004). Epistemic policies have the objective of increasing the accuracy of factual findings, or increasing the number of accurate findings, as well as improving the evidentiary warrant for findings and decisions. An example of an epistemic policy is allocating the burden of producing evidence to the party that is in the best position to produce that type of evidence. Examples of non-epistemic policy objectives are administrative efficiency and fairness to the parties. The reasoning that justifies a particular rule change should balance all of the epistemic and non-epistemic policies that are relevant to the proposed rule.

In law, as in ordinary life, the best new rules generally evolve from successful decisions in past cases. Clusters of factually similar cases can present a potential exception to be made within the existing legal rules. One task for the D-L Framework is therefore to analyze the patterns within the schematized evidence of cases. The D-L Framework allows

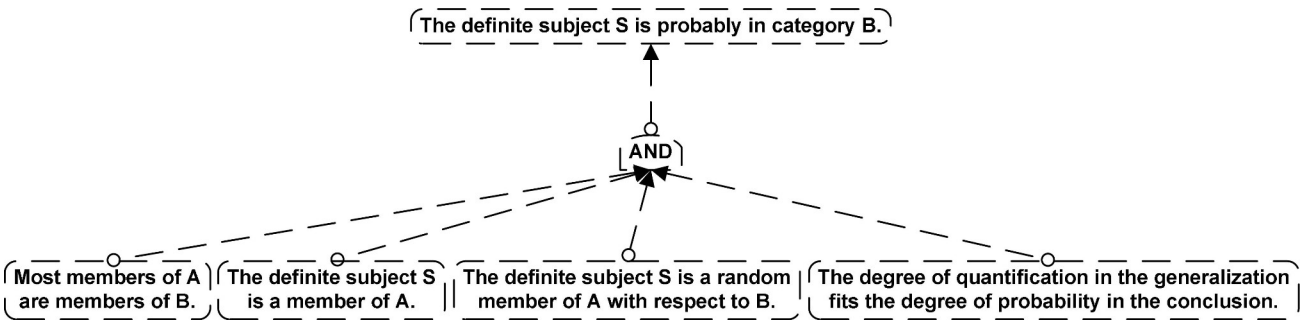

Figure 4. A D-L Framework Logic Diagram for the Direct-Inference Plausibility Schema

empirical study of how the law's conservative and incremental approach operates in actual cases.

A current hypothesis is that rule evolution in law involves at least two types of reasoning. The first uses higher-order rules (such as minimum due process) to constrain what kinds of rules are acceptable. The D-L Framework could model such process rules using implication trees. The second type of reasoning uses higher-order schemas to balance competing policies and objectives – for example, when there must be a trade-off between factfinding accuracy and administrative efficiency. The D-L Framework could model such policy-balancing reasoning by developing counterparts to plausibility schemas.

Evolving Plausibility Schemas

Plausibility schemas are designed to warrant default inferences to defeasible yet presumptively true conclusions. A major strategy for designing a plausibility schema is to develop a "*theory of uncertainty*" for the type of inference involved (Walker 2001). A theory of uncertainty explains how the available evidence could be plausible but the conclusion could be false (or in the case of a plausible defeater, how the conclusion could still be true). It identifies the possible sources of error inherent in the type of inference, and analyzes the sources, types, and degrees of uncertainty associated with drawing the conclusion. In examining the inherent uncertainty, however, a theory of uncertainty also explains why it is reasonable to draw the conclusion in a tentative way, even on the basis of incomplete evidence. Every plausibility schema, therefore, reflects a theory of uncertainty about why the schema's inference is defeasible yet warranted.

Theories of uncertainty are therefore one method of evolving new plausibility schemas. The advantage of the D-L Framework is that it can capture in a standard logical format the detailed reasoning of expert witnesses or other decisional agents in actual legal cases. Thus, the D-L Framework could assist the evolutionary process by identifying the patterns of factual reasoning that actually occur in those cases. Moreover, empirical research on that reasoning may suggest how to automate particular processes within the evolution of new plausibility schemas.

Coordinating the Application of Knowledge Models to Cases

The traditional legal distinction between rules of procedure and rules of evidence remains a useful distinction for process rules. *Procedural process rules* govern the dynamics and timing of default reasoning by authorizing particular procedural decisions under certain conditions. For example, early in a civil proceeding a defendant may move to dismiss the case for lack of jurisdiction, or any party may move for summary judgment before trial based on depositions and affidavits, or may move for directed verdict during trial based upon the testimony (Federal Rules of Civil Procedure; James, Hazard, and Leubsdorf 1992). Implication trees can model the rules for making such procedural decisions, and plausibility schemas can organize the relevant evidence in a particular case.

Evidentiary process rules govern the process of evaluating evidence and making findings about terminal propositions. Evidentiary decisions might apply rules about relevancy (attaching evidentiary assertions to terminal propositions); rules about admissibility (excluding some relevant evidence from the case altogether, or prohibiting its attachment to certain terminal propositions); rules about sufficiency of evidence (deciding whether the attached evidence can warrant a finding that a terminal proposition is true); standards of proof (establishing the threshold degree of plausibility required to find a terminal proposition to be true); and rules allocating the burden of persuasion (determining what finding to make when the attached evidence evaluates precisely on the threshold required by the standard of proof). An example of a particular evidentiary rule is admitting an expert opinion into evidence only if it is "scientific, technical, or other specialized knowledge" and it "will assist the trier of fact" (Federal Rules of Evidence; Daubert v. Merrell Dow Pharmaceuticals, Inc.; General Electric Co. v. Joiner; Kumho Tire Co. Ltd. v. Carmichael). Implication trees can model such evidentiary rules, and plausibility schemas can apply them in particular cases.

The D-L Framework integrates substantive and process reasoning by connecting process implication trees to the main implication tree for a legal decision. For example, a jurisdictional implication tree would be a high-level conjunctive branch for any implication tree for a court judgment. An evidentiary implication tree, on the other hand, might be a defeater branch connected near a terminal proposition of the main implication tree. When the same process implication tree may connect to various decisional trees, or may connect to many branches in the same decisional tree, it is efficient to model the process rules as a separate tree and connect that tree only as needed. The DA software, for example, can launch process trees from any point in any inference tree.

Conclusion

The domain of law provides a strategic area for studying multiagent, problem-oriented systems. Legal cases are numerous, complex, and socially important; legal reasoning is extensive and well-documented; legal decision-making integrates rules and policies with expert and non-expert evidence. The default-logic framework discussed in this paper successfully captures legal knowledge, integrates and evaluates expert and non-expert evidence, coordinates agents working on different legal problems, and evolves the knowledge model over time. A complete inference tree for the reasoning in a particular legal case consists of an implication tree that models all of the applicable substantive and process rules, together with the schematized evidentiary assertions attached to the terminal propositions of that implication tree. The process rules help to coordinate multiple agents and to evolve new implication trees and

plausibility schemas. The syntax and semantics of this framework allow the automation of key tasks, and the emergence of dynamic structures for integrating human and non-human agents. The logical basis of the framework ensures its applicability to knowledge and problem domains of similar complexity to law.

Acknowledgments

The author wishes to thank Hofstra University for its research support in preparing this paper.

References

American Law Institute. 1965. *Restatement of Torts, Second*.

Brewka, G., Dix, J., and Konolige, K. 1997. *Nonmonotonic Reasoning: An Overview*. Stanford, CA: CSLI Publications.

Chierchia, G., and McConnell-Ginet, S. 2000. *Meaning and Grammar: An Introduction to Semantics*. Cambridge, Mass.: MIT Press.

Copi, I. M., and Cohen, C. 1998. *Introduction to Logic*. Upper Saddle River, NJ: Prentice Hall.

Daubert v. Merrell Dow Pharmaceuticals, Inc., 509 U.S. 579 (1993).

Director, Office of Workers' Compensation Programs, Department of Labor v. Greenwich Collieries, 512 U.S. 267 (1994).

Dobbs, D. R. 2000. *The Law of Torts*. St. Paul, Minn.: West.

Federal Rules of Civil Procedure. 2006.

Federal Rules of Evidence. 2006.

General Electric Co. v. Joiner, 522 U.S. 136 (1997).

Gottwald, S. 2001. *A Treatise on Many-Valued Logics*. Baldock, England: Research Studies Press.

James, Jr., F., Hazard, Jr., G. C., and Leubsdorf, J. 1992. *Civil Procedure*. Boston, Toronto, London: Little, Brown.

Josephson, J. R., and Tanner, M. C. 1996. Conceptual Analysis of Abduction. In Josephson, J. R., and Josephson, S. G., eds. *Abductive Inference*. Cambridge, UK: Cambridge University Press.

Kadane, J. B., and Schum, D. A. 1996. *A Probabilistic Analysis of the Sacco and Vanzetti Evidence*. New York, NY: John Wiley & Sons.

Kumho Tire Co. Ltd. v. Carmichael, 526 U.S. 137 (1999).

Kyburg, Jr., H. E. 1983. The Reference Class. *Philosophy of Science* 50:374.

Kyburg, Jr., H. E. 1990. *Science & Reason*. New York, NY: Oxford University Press.

Kyburg, Jr., H. E., and Teng, C. M. 2001. *Uncertain Inference*. Cambridge, UK: Cambridge University Press.

Larson, R., and Segal, G. 1995. *Knowledge of Meaning: An Introduction to Semantic Theory*. Cambridge, Mass.: MIT Press.

Levi, I. 1977. Direct Inference. *The Journal of Philosophy* 74:5.

Levi, I. 1981. Direct Inference and Confirmational Conditionalization. *Philosophy of Science* 48:532.

Levi, I. 1996. *For the Sake of Argument*. Cambridge, UK: Cambridge University Press.

Malinowski, G. 1993. *Many-Valued Logics*. Oxford, UK: Oxford University Press.

Pollock, J. L. 1990. *Nomic Probability and the Foundations of Induction*. New York, NY: Oxford University Press.

Prakken, H. 1997. *Logical Tools for Modelling Legal Argument*. Dordrecht, The Netherlands: Kluwer.

Rodes, Jr., R. E., and Pospesel, H. 1997. *Premises and Conclusions: Symbolic Logic for Legal Analysis*. Upper Saddle River, NJ: Prentice Hall.

Russell, S. J., and Norvig, P. 2003. *Artificial Intelligence: A Modern Approach*. Upper Saddle River, NJ: Prentice Hall.

Saeed, J. I. 2003. *Semantics*. Malden, Mass.: Blackwell.

Sainsbury, M. 1991. *Logic Forms: An Introduction to Philosophical Logic*. Oxford, UK: Basil Blackwell.

Salmon, W. C. 1973. *Logic*. Englewood Cliffs, NJ: Prentice-Hall.

Schum, D. A. 1994. *Evidential Foundations of Probabilistic Reasoning*. New York, NY: John Wiley & Sons.

Singh, M. P., Rao, A. S., and Georgeff, M. P. 1999. Formal Methods in DAI: Logic-Based Representation and Reasoning. In Weiss, G., ed., *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*, 331-376. Cambridge, Mass.: MIT Press.

Toulmin, S. 1958. *The Uses of Argument*. Cambridge, UK: Cambridge University Press.

Toulmin, S., Rieke, R., and Janik, Allan. 1984. *An Introduction to Reasoning*. New York, NY: Macmillan.

Walker, V. R. 1996. Preponderance, Probability, and Warranted Factfinding. *Brooklyn Law Review* 62:1075-1136.

Walker, V. R. 1999. Language, Meaning, and Warrant: An Essay on the Use of Bayesian Probability Systems in Legal Factfinding. *Jurimetrics* 39:391-430.

Walker, V. R. 2001. Theories of Uncertainty: Explaining the Possible Sources of Error in Inferences. *Cardozo Law Review* 22:1523-70.

Walker, V. R. 2003. Epistemic and Non-epistemic Aspects of the Factfinding Process in Law. *APA Newsletter* 3,1:132-136.

Walker, V. R. 2004. Restoring the Individual Plaintiff to Tort Law by Rejecting "Junk Logic" About Specific Causation. *Alabama Law Review* 56:381-481.

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

Walton, D. 2002. *Legal Argumentation and Evidence*. University Park, Penn.: Pennsylvania State University Press.