Practical Reasoning Through Argumentation

Arthur M. Farley
Kathleen Freeman

CIS-TR-94-26
December 1994

Department of Computer and Information Science
University of Oregon

Practical Reasoning Through Argumentation

Arthur M. Farley
(art@cs.uoregon.edu)
Kathleen Freeman
(kfreeman@cs.uoregon.edu)

Computer and Information Science
University of Oregon
Eugene, OR 97403 USA

ABSTRACT

Practical reasoning is the process of drawing conclusions relative to evidence, goals, and risks. We present a computational model of dialectical argumentation as basis for practical reasoning in weak-theory domains, i.e., incomplete, uncertain, and inconsistent contexts. As information structure, an argument consists of argument units connecting claims with supporting data. As dialectical process, an argument consists of an alternating series of moves by opposing sides. Burden of proof, i.e., which side of an argument must realize what level of support, acts as move filter and termination criterion. Other aspects of our model include: (i) arguments defeat other arguments, not claims; (ii) not all conflicting arguments are defeating arguments; and (iii) whether a claim is believed or not depends in part upon the burden of proof specified.

Keywords: reasoning under uncertainty, knowledge representation, argumentation, burden of proof

INTRODUCTION

We present progress from research investigating the definition of a computational model of dialectical argumentation. The work is motivated by a desire to incorporate an ability to construct arguments into programs that attempt to generate intelligent behavior. We investigate the use of dialectical argument as basis for automated decisionmaking in weak theory domains (Porter, et.al, 1990), i.e., contexts wherein knowledge is incomplete, uncertain, or inconsistent. In such domains, a reasoner usually cannot derive claims that are deductively valid. As most claims are not established conclusively, a new definition of "proof" is needed. This problem is not new. Proposed solutions primarily have revolved around propagation of varying degrees of belief (Shafer and Pearl, 1990) or have employed notions of default reasoning and nonmonotonic logic (Ginsberg, 1987).

Being able to assign truth values other than 0 (false) or 1 (true) to propositions would appear to resolve issues of uncertainty by directly representing hedges on belief. However, this approach raises other issues regarding appropriate propagation rules and, more importantly, what is actually believed after propagation has been completed. Furthermore, computation of maximum a posteriori probabilities for all but the simplest belief network structures has been shown to be NP-complete (Shimony, 1994).

Non-monotonic and default reasoning systems have the property of being able to entertain assumptions regarding the truth of propositions and to realign truth valuations should further information or decisions contradict these assumptions. The unresolved question of what is the smallest adjustment to a set of beliefs or what to believe under conditions of multiple extensions (i.e., several consistent subsets of propositions) again makes this solution to belief assignment in weak theory domains incomplete.

To resolve the question of ultimate belief, left hanging by the above solutions, the notions of *sceptical* and *credulous* reasoners have been defined (Touretzky et.al., 1987). Sceptical reasoners withhold belief when encountering (equally) good reasons for belief in a conclusion and its negation. Credulous reasoners make choices between controversial conclusions at random, e.g., by flipping a coin, so as to get off the dime in contradictory, confusing circumstances. These choices represent but two of the various possible positions along a scale of credulity, but remain largely unresponsive to the need of practical reasoning systems to be able to make appropriate, not random or no, decisions in weak theory domains.

Argumentation, with its emphasis on both supporting and refuting claims under situations of uncertainty and inconsistency, is well suited to serve as a framework for a practical definition of proof and proof procedure. Furthermore, burden of proof, loosely defined as which side of an argument must realize what level of support, becomes an important aspect of this framework, allowing us to define various, practical notions of proof. One context in which the notion of burden of proof has been defined historically and applied formally is the legal domain. Different burdens of proof are mandated at different stages of the legal process and for different types of legal action. These notions have been introduced to avoid sceptical judges inability to make reasonable choices and credulous judges flipping coins. We will provide operational definitions for several burdens of proof that are inspired by those used in legal settings. These will provide us with means to position a practical reasoning system at appropriate locations between sceptical and credulous reasoners as the situation demands.

MODELING ARGUMENT STRUCTURE

The representation of an argument as a structured entity and the generation of an argument as a dialectical process are both crucial elements of our theory. For argument as supporting explanation, we create argument structures that serve to organize relevant, available, and plausible support for a claim and its negation. We represent these argument structures in a form derived from that described by Toulmin in <u>The Uses of Argument</u> (1958). An argument comprises *data* (i.e., input evidence, grounds) supporting or refuting a *claim*. The basis for a connection between data and claim or the authorization for moving from data to claim is called a *warrant*.. Data and warrant may not be enough to establish a claim conclusively; as such, a claim has *qualifications*. Furthermore, the claim may be subject to *rebuttals*, i.e., circumstances where its support would not hold or where its negation is supported.

In our representation of argument, all claims, including input data, must be somehow supported, i.e., have *backing*. We define two types of backing: *atomic*, for information from outside the immediate realm of the argument, and *tau* ("<u>T</u>oulmin <u>a</u>rgument <u>u</u>nit"), where the claim is supported by data through application of a warrant. A claim may have multiple backings. A rebuttal is a rival claim (i.e., the negation of a claim) and the arguments that support that claim. A warrant is a rule-like piece of knowledge, having antecedent and consequent aspects. The antecedent and consequent fields consist of one or more propositional clauses. Multiple clauses in either the antecedent or consequent are taken to represent conjunctive elements.

A warrant also has two type fields. The *wtype1* field classifies the relationship between the antecedent and consequent as explanatory (ex) or sign (si), as in (Freeley, 1990). An example of an explanatory relationship is a causal link, where knowledge of the antecedent "explains" knowledge of the consequent, e.g.,"fire" causes (explains) its consequent "smoke". A sign relationship represents a correlational link between antecedent and consequent, e.g., "Summer weekends are generally rainy." The *wtype2* field of a warrant represents the strength with which its consequent can be drawn from its antecedent. Possible values are sufficient (s), default (df), and evidential (ev). A sufficient warrant is meant to represent conclusive relationships, e.g., definitions. Default and evidential warrants are meant to represent two levels of uncertain knowledge, with default indicating relationships that are usually (almost always) the case (e.g., "birds fly") and evidential referring to less certain, but still likely, links (e.g., "persons who live in Bermuda are more often British subjects"). Warrants are expected to be written in the direction that accommodates the strongest possible type.

TABLE 1. Reasoning Steps

warrant	data	conclusion	reasoning step
p -> q	p	q	modus ponens (MP)
p -> q	not q	not p	modus tollens (MT)
p -> q	q	p	direct abduction (ABD)
p -> q	not p	not q	contrapositive abduction (ABC)

Given a warrant with antecedent p and consequent q, we define allowable reasoning steps in Table 1. The latter two reasoning steps are fallacies in the context of deductive reasoning (asserting the consequent and denying the antecedent, respectively). However, they can be appropriate for reasoning when knowledge is incomplete or uncertain. Polya (1968) discusses such "patterns of plausible inference", calling them "examining a ground" (MP, ABC) and "examining a consequent" (MT, ABD).

When deductive and plausible reasoning types are present in the same system, care must be taken to avoid inappropriate reasoning combinations (Pearl, 1987). For example, if the reasoner knows that "rain causes wet-grass" and "sprinkler-on causes wet-grass", an unrestricted combination of modus ponens and direct abductive reasoning would allow the reasoner to derive support for the claim "sprinkler-on" from the input data "rain". To block the generation of such unacceptable arguments, MP/ABD combinations are not permitted over two explanatory warrants, demonstrating the need for our wtype1 distinction on warrants.

Qualifications are used to capture the level of support for a claim, reached as a result of arguments based upon uncertain knowledge and plausible reasoning. Presently, we use the following qualifications: valid(!), strong (!-), credible (+), weak(-), and unknown (?). The first four are ranked in order of decreasing level of support; while the last indicates no support in the current argument. The qualification on a claim is that associated with its strongest supporting argument. The qualification on a claim with tau backing is the least of the qualifications associated with the warrant application, being the qualification(s) on the data support, on the warrant itself, and that derived from the warrant type and reasoning step applied ("link qualification"; see Table 2).

This weakest link approach to propagating support for argument claims and its appropriateness for plausible reasoning has been discussed elsewhere (Pollock, 1991; Rescher, 1976). We attempt to appropriately capture modus tollens reasoning, by representing it as a weak reasoning method when not involving a sufficient warrant.

Each claim has two associated qualifications, summarizing the argument strength for the claim and for its negation; as such, we represent all claims in an argument structure in positive (i.e., unnegated) form.

TABLE 2. Link Qualifications

warrant type	reasoning step	link qualification
$\to_s$	MP, MT	valid
$\to_s$	ABD, ABC	weak
$\to_{df}$	MP	strong
$\to_{df}$	MT, ABD, ABC	weak
$\to_{ev}$	MP	credible
$\to_{ev}$	MT, ABD, ABC	weak

MODELING ARGUMENT GENERATION

Having only a structural model of argument does not capture the procedural, sequential character of dialectal argumentation. Dialectical argumentation results in the intertwining, over time, of argument structures generated by Side-1 in support of a claim and those from Side-2 that support its negation. Support for an input claim, we will term as *Side-1 support*; support for its negation we term *Side-2 support*.

Dialectical argumentation begins with Side-1 attempting to find support for the input claim. Finding support for a claim results in the generation of argument structures as described earlier. Given a claim, search for support proceeds in a backward-directed fashion from the input claim toward input data, using warrants as intermediate steps. The process grounds out when a (sub)claim is supported by a proposition in the input. A new tau structure is generated for each warrant applied; the qualification and backing fields of the claims are updated to reflect the new support. If no initial support can be found, the argument ends with a loss for Side-1; all burdens of proof require that at least one supportive argument for the input claim be found.

If Side-1 is able to find support for the claim, control passes to Side-2, which tries to refute the argument for claim(s) established by Side-1. We distinguish two types of refutation: (a) rebutting and (b) undercutting, in a manner related to Pollock (1987). *Rebutting* finds new arguments directly supporting the negation of a claim. *Undercutting* is accomplished by finding weaknesses in purported support for a claim, questioning the sufficiency of the input support or of tau fields (i.e., warrant type and reasoning type). We undercut support for one claim by rebutting the support for one of its supporting claims according to the structure of taus. Argument moves implementing these tasks of dialectical argumentation are described in Table 3.

If an undercutting move is successful, it may result in withdrawal of an argument. Such moves are said to be *defeating* arguments and are indicated by the * entries in Table 3. These moves are in response to an argument for which an exception is found or to a weak argument made by the other side, i.e., those based only on plausible, not deductive, reasoning steps. Rebutting arguments that merely find alternative, unrelated arguments for the negation of a claim only serve to make the original conclusion controversial. Whether this is sufficient for a given side of the argument will depend on the burden of proof required.

TABLE 3. Dialectical Argument Moves

ARG TASKS	MOVES	GIVEN	SHOW	defeat
support C	(a) support	C	X->C ^ X	
			C->X ^ X	
			~C->X ^ ~X	
			X ->~C ^ ~X	
refute C				
undercut C	(b) invalid	X->C	~X	
	antecedent	^X		
	(c) exception	X->C	X^Y->~C	*
		^X	^Y	
	(d) inapplicable	X->~C	Y->~C	*
	evidence	^~X	^Y	
	(e) unneeded	C->X	Y->X	*
	explanation	^X	^Y	
rebut C	(f) reductio ad	C	C-->Z	
	absurdum		^~Z	
	(g) rival	C	X->~C	
	support		^X	
	(h) missing	C	X->C	
	support		^~X	
	(i) rival	C	~C->X	
	implication		^X	

Note that arguments defeat (steps in) other arguments, threby impacting support for claims; unrelated, rebutting arguments do not defeat arguments for a claim, but simply make the claim controversial. For example, suppose we put forth the argument for the claim that a penguin flies because it is a bird. An argument based on a warrant that most things whose name start with p don't fly would only make the claim controversial. However, our initial argument for penguins flying would be defeated by the argument that penguins are an exceptional sort of bird that does not fly.

When a side is in control of the argument, it must select which argument move(s) to apply. Heuristics that order argument moves for selection are meant to reflect two goals: generate the strongest arguments possible for the active side and generate coherent arguments, i.e., arguments that are responsive to those put forward by the other side. As such, agument moves are ordered, as follows: (a) valid reasoning steps are preferred over plausible steps; (b) moves that are defeating are preferred over moves that only make a claim controversial; (c) moves that attack a supporting argument closer to the overall claim are preferred; and (d) undercutting moves are preferred over rebutting moves. Warrants are also ordered according to the following criteria: (a) specific warrants (i.e., those with more antecedents) are preferred over more general warrants; (b) stronger warrant types are preferred; and (c) warrants for which the antecedent currently has consistent or no known input support are preferred.

These ordering heuristics anticipate moves the other side may use to try to refute a claim. Strong reasoning steps are more difficult to defeat; those closer to the root claim leave fewer opportunities for alternative support; defeating arguments eliminate controversial elements; weaker reasoning types allow more opportunities for defeating refutations. Controversial or negated data can be used to support a claim weakly at best.

This completes an overview of the basic elements of our model of argumentation. Given a set of warrants, some input data, a claim, and a burden of proof, our system proceeds to generate a dialectical argument. Control switches from side to side as check conditions, i.e., sufficient refutations, are realized. Deciding which moves are sufficient to generate a check condition for a particular side, when an argument process is complete, and who wins, all depend upon a given burden of proof.

BURDEN OF PROOF

Now we turn our attention to the definition of burden of proof and discuss its impact on argument generation and outcome. There are two elements to the notion of burden of proof: (1) which side of the claim bears the burden; (2) what level of support is required. As we consider only two sides of an argument (positive or negative), we assume that Side-1 always bears the burden of proof for the input claim, which could be stated as the negation of a proposition.

To understand the notion of burden of proof, suppose you are considering taking some action in a typical, real-world context. Examples of different burdens of proof you might request of your practical reasoning system are as follows: "If you can give me one good argument for this action, I will do it." "I need an iron-clad argument for this action before I will proceed." "Unless you can give me a predominating argument against this action, I will go ahead." These requests illustrate the two aspects of burden of proof: which side and what level. They also suggest a role for burden of proof in allocating risk. If one is not risk averse in a given situation, then one may just need a single, sound argument to proceed. If the cost of being wrong is high, then more strict requirements may be imposed. If an action is required to prevent disaster, one may be concerned only about whether there are any arguments against a proposed action. A *defendable argument* is one that can not be successfully defeated.

We define several levels of support, as follows:

- *scintilla of evidence (se)*
  find at least a weak, defendable argument
- *preponderance of the evidence (pe)*
  find at least a weak, defendable argument
  outweigh the other side's arguments
- *dialectical validity (dv)*
  find at least a credible, defendable argument
  defeat all of the other side's arguments
- *beyond a reasonable doubt (brd)*
  find at least a strong, defendable argument
  defeat all of the other side's arguments
- *beyond a doubt (bd)*
  find a valid, defendable argument
  defeat all of the other side's arguments

Burden of proof plays several roles in the process of argumentation: (i) as basis for deciding relevance of particular argument moves, (ii) as basis for deciding sufficiency of a side's move (i.e, whether a check condition has been realized); (iii) as basis for declaring an argument over; and (iv) as basis for determining the outcome (i.e., belief decision) of an argument. For example, if we have imposed a burden of proof of dialectical validity and Side-2 has presented an argument refuting Side-1's claim, Side-1 can not merely find another argument supporting the input claim; Side-1 must defeat the refutation or admit defeat. However, if the burden of proof is only preponderance of the evidence, then another argument in favor of the claim by Side-1 may be sufficient to outweigh Side-2's rebuttal. For a burden of proof of beyond a reasonable doubt, unless Side-1 can find an initial argument based upon valid application of a sufficient or default warrant, it must concede defeat without Side-2 even needing to make a move as no strong support can be found.

EXAMPLES

We demonstrate our notions of argumentation and burden of proof as basis for practical reasoning by considering two examples that have been previously discussed in the literature. First, we discuss variants of the problem known as the "Nixon diamond" (Poole, 1989). The knowledge for the argument is represented by the following warrants and data:

```
(w1 ((republican))  --> si ev ((hawk)) (!? GIVEN))
(w2 ((quaker))  --> ex df ((dove)) (!? GIVEN))
(w3 ((hawk)) --> ex df  ((supports star wars)) (!? GIVEN))
(w4 ((hawk)) --> si df ((politically motivated)) (!? GIVEN))
(w5 ((dove)) --> si df ((politically motivated)) (!? GIVEN))
(w6 ((quaker)) --> ex df ((religious)) (!? GIVEN))
(w7 ((hawk) --> ex s ((not (dove)) (!? GIVEN))

(d1 (republican) (!? GIVEN))
(d2 (quaker) (!? GIVEN))
(claim (dove) (?? NIL))
```

That is, being a Republican is generally correlated with being a hawk, while someone who is a Quaker is almost always a dove. Someone who is a hawk usually supports star wars. Both hawks and doves are usually politically motivated; being a Quaker explains being religious in most cases. Finally, a hawk is not a dove, by definition. The initial data consists of valid information that Nixon is a Republican and a Quaker. The input claim is that Nixon is a dove.

With a burden of proof of scintilla of evidence, Side-1 is able to find strong support for the input claim (dove) through MP application of warrant w2 based on input data d2. Side-2 is unable to defeat this argument; thus, under a scintilla of evidence, (dove) is believed. Note that if the input claim were the opposite, i.e., (not (dove)), then MP application of w1 based on input d1 followed by MP application of w7 would provide a credible, undefeatable argument for that side of the claim, as well. Thus, both (dove) and (not (dove)) would be believed with a scintilla of evidence.

If we strengthen the required proof level to preponderance of evidence, Side-2 has to find support that matches the strong support of Side-1's original argument given above. There is no opportunity for this with the given knowledge; while a credible argument for (not (dove)) exists, as discussed above, no strong argument exists. The claim (dove) is believed with burden of proof of preponderance of the evidence. If (not (dove)) were the input claim, then Side-2 can counter the credible support for the input claim with strong support for (dove), and Side-1 would lose the argument. Burden of proof levels of dialectical validity and above require that Side-1 directly defeat any counterargument provided by Side-2. At these levels, we can establish neither of the claims (dove) or (not(dove)). There are no opportunities for defeating arguments related to these claims in the given knowledge; only controversy can be raised.

If we consider (supports-star-wars) as the input claim, we see that the only way to argue for (not (supports-star-wars)) is through ABC application of warrant w3, based upon the strong argument for (not (hawk)) from MT application of w7 following MP application of w2 based upon input (quaker). This argument can be directly defeated by Side-2 through MP application of w3, based upon an argument for hawk using warrant w1 and input (republican). Thus, (supports-star-wars) would be believed with a burden of proof of dialectical validity; as such, we see that the claim (not (supports-star-wars)) loses all arguments. Similarly, we would conclude (religious) and (politically-motivated) for burden of proof up through beyond a reasonable doubt. The only arguments against these claims are weak and easily defeated by conterarguments.

As a final exercise, consider that there is only credible (e.g., hearsay) evidence that Nixon is a Quaker; the qualification on input d2 becomes (d2 (quaker) (+? GIVEN)). In this case, at most a credible argument can be realized for both (dove) and (not (dove)). Thus, neither claim is able to win a preponderance of the evidence argument, indicating that the given knowledge is even more highly inconclusive in this case.

Poole (1990), in discussing multple extensions as derived for this problem, considers a series of prediction mechanisms: (1) predict only logical consequences of the input facts; (2) predict only what is incontestably explainable (could include logical consequences of hypotheses as well as facts); (3) predict what is in every extension; (4) predict what is in any extension; (5) predict what is not inconsistent with the known input facts. He says "It does not seem reasonable to be less sceptical that (5) or...be more sceptical than (1)".

This last is a point of departure between our two approaches, as we can in fact be less skeptical than (5). At some levels of burden of proof, we allow conclusions to be drawn that are inconsistent with other conclusions. Though this appears illogical, it is not impractical (irrational or unreasonable). For example, in the case of the Nixon diamond, it is the facts themselves (quaker/republican) that are really at the heart of the controversy regarding dove. Thus, to require that our conclusions be consistent with what is "known" is too strong a constraint to hope for or limit ourselves to. In viewing our model of argumentation as practical reasoning, we provide preponderance of evidence and scintilla of evidence as weak burdens of proof in favor of a claim or allow for demand of high level of proof for the claim's negation as ways to weaken the strength of arguments required for belief. If one does want to enforce stronger assurances on beliefs, proof levels at or above dialectical validity are available, also.

The second example we consider has been termed the "lottery paradox" (Pollock, 1994). In rough translation, suppose there is a fair lottery consisting of a large set of numbered tickets, from which one and only one sold ticket will be drawn as the winner. The practical question is, "should one buy a ticket?" Given that one's interest is only in winning the lottery, not helping the worthy cause running it, the practical question is "will one win the lottery, if one buys a ticket?"

We represent the problem situation by the following warrants and input data:

```
(w1 ((wins-lottery)) => ex df
                  ((fair-lottery) (buys-ticket) (ticket-drawn) (!? GIVEN))
(w2 ((fair-lottery) (buys-ticket) (ticket-drawn)) => ex s
                                            ((wins-lottery)) (!? GIVEN))
(w3 ((not (buys-ticket)))  => ex df ((not (ticket-drawn))) (!? GIVEN))
```

(w4 ((fair-lottery) (buys-ticket)) => si df ((not (wins-lottery))) (!? GIVEN))

(d1 (fair-lottery) (!? GIVEN)
(d2 (buys-ticket) (!? ASSUMED)
(claim (wins-lottery) (?? NIL))

The warrant w1 captures the notion that if one wins the lottery, usually it was a fair lottery, one bought a ticket and the ticket was drawn (of course, the lottery might have been unfair, so only a default warrant). Warrant w2 captures the opposite notion that, by definition of fair lottery, if one buys a ticket and the ticket is drawn then one wins the lottery. Warrant w3 captures the notion that if one doesn't buy a ticket then one's ticket can't be drawn (again, with a slight chance that lottery is unfair). Finally, warrant w4 represents the unlikelihood of winning such a lottery if one enters.

The least demanding burden of proof would be scintilla of evidence for the input claim (wins-lottery). Side-1 could choose to apply warrant w1 with the abductive reasoning step ABD or warrant w2 with reasoning step MP; by our ordering principles, warrant w2 is chosen. While (fair-lottery) and (buys-ticket) are part of the input, Side-1 must find tau support for (ticket-drawn). This is realized through warrant w3, applied in abductive form ABC, where (buys-ticket) weakly supports (ticket-drawn); by weakest link propagation of support, this yields weak support for the conclusion (wins-lottery). Side-2 can only defeat this argument by attacking the abductive step with w3. With the knowledge given, no such defeat is possible and (wins-lottery) is believed with scintilla of evidence. If the burden of proof were on the same side to realize preponderance of the evidence, Side-2 can now counter with MP application of warrant w4, providing strong support for (not (wins-lottery) and winning the overall argument.

If we made the input claim (not (wins-lottery)), demanding proof one couldn't win before not buying a ticket, we could set the proof level to scintilla of evidence indicating high risk aversion. If this case Side-1 uses warrant w4 to provide a strong argument for not winning. Side-2 can potentially defeat that argument with an argument for an exception through warrants w2 and w3, as above. However, the weak strength of this argument is insufficient to defeat the argument given by Side-1. A potentially defeating argument must have strength at least equal to that of the defeated argument to be a defeater. We saw before that (not (wins-lottery)) also wins preponderance of the evidence. However, at proof levels above that, (not (wins-lottery)) loses, as while only a weak argument for winning can be had through warrants w2 and w3 it can not be defeated.

For the lottery paradox, we find that in cases where only scintilla of evidence is required for it or dialectically valid is required against it, (wins-lottery) wins the argument. Thus, there are a number of situations under which a practical reasoning system based on argumentation would choose to buy a ticket to win the lottery, essentially when risk aversion is low. This is in contrast to other systems that can only decide on such a course of action by flipping a coin after adopting credulous reasoning.

CONCLUSION

We see that burden of proof is a particularly useful aspect of a computational model of argumentation as basis for practical reasoning. Our model comprises both argument as supporting explanation and argument as dialectical process. It incorporates other features appropriate for reasoning in weak theory domains, including plausible inference and uncertainty representation. We demonstrate the application of different burdens of proof in two simple argument contexts

Our model of dialectical argumentation has been implemented and evaluated on a number of classic reasoning problems in weak theory domains. The implemented model exhibits reasonable behavior when applied to these benchmark examples from formal argumentation and artificial intelligence (Freeman, 1993). We hope this model can serve as a framework for further exploration of argumentation as a means for practical reasoning. We are investigating several extensions to our model, including addition of an "anecdotal" warrant type to incorporate case-based reasoning and inclusion of other means for uncertainty representation.

REFERENCES

Freeley, A. (1990). Argumentation and debate: Critical thinking for reasoned decision making (7th ed.). Belmont, CA: Wadsworth Publishing Company.

Freeman, K. (1993). Toward Formalizing Dialectical Argumentation. PhD Dissertation, Department of Computer and Information Science, University or Oregon.

Ginsberg, M.L. ed. (1987) Readings in Nonmonotonic Reasoning, Los Altos, CA, Morgan Kaufmann.

Pearl, J. (1987). Embracing causality in formal reasoning. Proceedings of the Sixth National Conference on Artificial Intelligence, AAA1-87, 369-373.

Pollock, J. (1991). A theory of defeasible reasoning. International Journal of Intelligent Systems, *6*, 33-54.

Pollock, J. (1992). How to reason defeasibly. Artificial Intelligence, *57*, 1-42.

Pollock, J. (1994). Justification and defeat. Artificial Intelligence, *67*, 377-407.

Polya, G. (1968). Mathematics and plausible reasoning (2nd ed.) (vol. II). Princeton, NJ: Princeton University Press.

Porter, B., Bareiss, R., & Holte, R. (1990). Concept learning and heuristic classification in weak theory domains. Artificial Intelligence, *45*, 229-263.

Rescher, N. (1976). Plausible Reasoning. Assen/Amsterdam, The Netherlands: Van Gorcum.

Shafer, G. and Pearl, J. eds. (1990), Readings in Uncertain Reasoning, Los Altos, CA, Morgann Kaufmann.

Shimony, S.E. (1994) Finding MAPs for belief networks is NP-hard. Artificial Intelligence, *68*, 399-410.

Toulmin, S. (1958). The uses of argument. Cambridge: Cambridge University Press, 1958.

Touretzky, D.S., Horty, J.F., and Thomason, R.H. (1987) A clash of intuitions: The current state of nonmonotonic multiple inheritance systems, Proceedings of IJCAI-87, Milan, It., 476-482.