Argumentation in Weak Theory Domains

Kathleen Freeman & Arthur Farley

CIS-TR-92-10
April 1992

Department of Computer and Infomation Science
University of Oregon

Argumentation in Weak Theory Domains

Kathleen Freeman
Department of Computer and Information Science
University of Oregon
Eugene, Oregon 97403
503-346-1381
503-346-5373 (fax)
kath@cs.uoregon.edu

Arthur M. Farley
Department of Computer and Information Science
University of Oregon
Eugene, Oregon 97403
503-346-4408
503-346-5373 (fax)
art@cs.uoregon.edu

Abstract

Reasoning depends on a reasoner's knowledge and beliefs. Yet, in real world domains, what is known or believed is often incomplete, inconsistent, or otherwise uncertain. We call such domains "weak theory domains".

We explore the use of argumentation as a basis for reasoning in weak theory domains. Argumentation is a method for locating, highlighting, and organizing relevant information in support of and counter to a plausible claim. This information can then serve as a vehicle for comparing the merits of competing claims.

We present aspects of our preliminary investigation of a formal theory of argumentation: (i) identifying a formal theory of argumentation; (ii) implementing the theory in a computer program; (iii) gathering example problems and associated arguments; and (iv) evaluating the theory with respect to the example arguments.

Current work concentrates on the structure and generation of independent, supporting arguments for an input claim and its negation. Future work will focus on argumentation as a series of adversarial moves that support and counter an input claim.

I. Introduction

The ability to reason is a hallmark of intelligence and therefore an essential skill to be realized by artificial intelligence systems. Reasoning (i.e., making inferences, solving problems, revising beliefs) depends on a reasoner's current knowledge and beliefs. Yet, when we turn our attention to real-world domains, we find that what is known or believed is often incomplete, ambiguous, imprecise, and inconsistent. In such domains, there will always be gaps in what is known due to a lack of information gathering capabilities, the combinatorial complexity of the space of possible situations, the possibility of exceptions to general rules, and the prevalence of changing circumstances.

We call real-world domains about which knowledge is incomplete, inconsistent, or otherwise uncertain, "weak theory domains". Expert systems "tend to be built precisely for domains in which decisions are highly judgmental" [2] and are thus expected to operate in weak theory domains. Various techniques for belief propagation and approaches to fault diagnosis have been developed in an attempt to cope with the uncertainty and unreliability in these domains (see, e.g., [5],[12]). Many of the problems studied in artificial intelligence research on knowledge representation and automated reasoning reflect the consequences of dealing with weak theory domains.

As an example, consider the following "Bermuda" problem, based upon the classic example of Toulmin [26]:

We know that most people with British passports are British citizens and that usually anyone born in Bermuda can be assumed to be a British citizen. However, someone born in Bermuda to alien parents is not a British citizen. Most people that are English speaking and have a Bermuda social security card were born in Bermuda. Finally, that a person has a Bermuda social security card is sufficient evidence to conclude that the person has Bermuda working papers. We have just been introduced to Harry, who speaks English, has a German passport, and shows us his Bermuda working papers. We are wondering whether he is a British citizen or not.

The knowledge in this problem is inconclusive, simultaneously supporting both conflicting conclusions. It is incomplete, since information that could help support one conclusion or the other, such as whether or not Harry's parents were aliens, is missing. Also, the knowledge given about the problem is uncertain. Hedges such as "often" and "usually" are indications that a rule is less than certain.

There are other sources of uncertainty, as well. Should the knowledge that Harry has working papers support the conclusion that he has a Bermuda social security number (and thereby support the born in Bermuda assertion)? Or, is the lack of a British passport sufficient grounds for claiming that Harry is not a British citizen? How much support does either of these give to the conclusion? How should uncertainty be propagated along a chain of reasoning, e.g., from working papers to social security card to Bermuda-born and finally to British citizen? How should support be combined across the various arguments that support the Harry's being a British citizen or not being a British citizen?

As we see above, support for a claim that is accumulated by reasoning in a weak theory domain may be inconclusive or controversial. Claims supported by *plausible reasoning* (e.g., [18]) based upon uncertain, incomplete knowledge are said to be *defeasible* (e.g., [17]) due to the uncertainty of generalizations and the availability of arguments to the contrary.

If we can not decide the truth of most claims with certainty, how can we proceed with reasoning at all? Our approach is to explore the use of argumentation as a basis for decision-making in weak theory domains. A reasoner should be able to justify claims made, both on their supporting merits and with respect to support for alternative claims. Weak links in the support for a claim should be explicated; links that can be countered by support for alternative claims should be discounted. Finally, there is a need for methods to compare the strength of support among alternative claims.

Our proposal is that argumentation be used as a method for locating, highlighting, and organizing relevant information in support of and counter to a plausible claim. Argumentation can be viewed as a mechanism for generating a class of super-explanations: in addition to providing "logical" support for belief in a claim (i.e., as we would expect to find in a typical explanation), reasoning that counters the claim is determined, as well. Argumentation is seen to be a vehicle for comparing the merits of support for competing claims. Such an ability is crucial for the justification of decision-making by artificial intelligence systems that hope to perform effectively in weak theory domains.

Argumentation has long been viewed as an important reasoning tool outside artificial intelligence (e.g., [22]). In artificial intelligence and elsewhere, argumentation has been variously investigated from the standpoints of rhetoric (e.g., [10],[11]), philosophy (e.g., [22],[26]), discourse analysis (e.g., [6]), legal reasoning (e.g., [1],[14],[15],[24]), and others. Our work borrows from this previous work, generalizing and refining it so as to focus on a *formal* theory of argumentation that can provide *decision support and justification* based upon *plausible* reasoning.

In the remainder of this paper, we present aspects of our preliminary investigation of argumentation. Our initial goals have been: (i) identify a formal theory of argumentation; (ii) implement the theory in a computer program; (iii) gather example problems and associated arguments that demonstrate various aspects of argumentation; (iv) evaluate the theory with respect to the example arguments. Our results will be of interest to researchers in both artificial intelligence (AI) and argumentation. For AI, the ability to generate arguments will be a useful technique in real-world contexts. For argumentation researchers, AI methodology offers a new way for formalizing and evaluating theories of argumentation.

II. A Model of Argument

In our approach to argumentation, we make two significant, simplifying assumptions: (i) of the five aspects of classical rhetoric (invention, arrangement, style, memory, and presentation/delivery), we concentrate on invention, i.e., the process of developing a defensible "proof" for an assertion; (ii) of the three types of "proof" in classical rhetoric (logos, ethos, pathos), we deal entirely with logos, i.e., logical reasoning (both informal and formal logic). These restrictions are intended to focus the work on the more formal aspects of argumentation.

Our approach is based on the following, complementary definitions of argument: (i) "the grounds . . . on which the merits of an assertion are to depend" [26], and (ii) "a method for conducting controversial discussions, with one contender defending a thesis in the face of object[ions] and counterarguments made by an adversary" [22]. Any claim is presumed to be controversial, i.e., in need of support and vulnerable to objections against which it must be defended. This presumption is appropriate for most claims made in weak theory domains.

There are two senses of argument indicated by the above definitions. The first defines argument as a supporting explanation, as in 'she made an argument'. The second concentrates on argument as an activity in which two or more agents engage. We see an argument as a method, or process, as well as a product. Thus, the generation and representation of arguments are both crucial to our theory.

Our theory will develop both senses of argument in successive stages. In the first stage, argument structures for a claim and its negation are generated. These two structures summarize the relevant, available knowledge about the claim, its support, and support for its negation. In this paper, we give a description of our theory of argument structure and its generation and presentation by a current implementation. Later, we outline approaches we intend to follow in the next stage, developing a model of the process of argumentation as interaction that will be based upon the argument structures described here.

We represent an argument in a modified version of the form suggested by Toulmin in his book *The Uses of Argument* [26]. According to Toulmin, an argument comprises input *data* (i.e., a problem description, evidence, grounds) said to support a *claim* (i.e., conclusion, solution, belief). The authority for taking the step from data to claim is called a *warrant*. The warrant may have *backing*, or justification. The data and the warrant may not be enough to establish the claim with certainty, i.e., the resultant claim may be *qualified* (e.g., "probably"). The claim may be subject to *rebuttals*, special circumstances where the warrant would not hold (i.e., "unless...").

Modifications to this structure are needed to (1) formalize Toulmin's ideas (especially warrants and qualifications); (2) provide a macro structure for arguments, e.g., extended chains of support for claims, multiple backings for claims; and (3) explicate various sources of uncertainty, i.e., arguable points in the domain knowledge.

An argument is represented as a structure of claims, interconnected by "Toulmin argument units" (*taus*). A tau is essentially the data-warrant-claim structure described above. But data and warrants are also viewed as claims, and an argument structure is a hierarchic structure of claims and taus, with taus supporting claims as conclusions and depending on claims as support. Since all the major elements of a tau are claims, we will refer to the tau as a data-warrant-*conclusion* structure to avoid ambiguity.

In addition, we make the following modifications/extensions to the basic Toulmin argument structure. Claims comprise data, or a statement of the claim, a qualification, and a backing. There are two types of backing: *atomic*, as input information from outside the domain of argumentation (i.e., inartistic proofs [10]), and *tau*, where the claim is supported by data and a warrant. A claim may have multiple supporting arguments (backings); the qualification on a claim summarizes the strengths and weaknesses of all the supporting arguments. In addition, a claim has a pointer (possibly nil) to its negation and, thus, implicitly to its rebuttals, which are the arguments supporting its negation.

Warrants have a slightly different structure from other claims. In addition to qualification, backing, and rebuttal, warrants have two data fields, *antecedent* and *consequent*. A warrant also has a *type* associated with it which represents information as to the strength of the connection between antecedent and consequent. While a warrant has antecedent and consequent, indicating the normal direction of its application, we can use warrants in the opposite direction as well. Given a warrant with antecedent p and consequent q, we define these *reasoning steps* as follows:

data	**conclusion**	**reasoning step**	
p	q	modus ponens	(MP)
~q	~p	modus tollens	(MT)
q	p	abduction	(ABD)
~p	~q	contrapositive abduction	(ABC)

The last two reasoning steps are fallacies (asserting the consequent and denying the antecedent, respectively) for deductive reasoning, but are often used in plausible reasoning to indicate support for a claim, though not conclusively. This leads to the need to attach qualifications to claims, as does the use of uncertain warrants. We next define a formal language for claim qualifications and warrant types. We then specify their interaction with the four types of reasoning steps indicated above in determining resultant claim qualifications.

III. Representing Uncertainty

As a first approximation to the representation of degrees of belief in a claim, we restrict qualification values to be one of the following: *sure, usual, probable, possible,* and *contingent*. These are ranked in decreasing order of degree of belief as given. While we are not adopting a strict probabilistic interpretation here, the first four can be related to probability values of 1, $> 1-\partial$, $> .50$, and > 0.0; contingent indicates a lack of input data support. This broad-scale quantification of degree of belief limits our ability to capture differences among strengths of some arguments. However, we later demonstrate instances where we can make important distinctions using this highly restricted set of values.

The qualification on a claim is that associated with its strongest supporting argument. The qualifications on (input) data are given as atomic backing at input time and remain unchanged thereafter, unless better support can be derived from tau backing. The strength of any claim drawn from application of a warrant (i.e., a tau backing) is the least of the qualifications associated with the application: the qualification on data serving as antecedents, the qualification on the warrant, and that derived from the type of warrant and reasoning step applied, as discussed below.

With each warrant we associate a type, reflecting the strength with which its conclusion can be drawn from the given antecedent. The types we specify are: *necessary and sufficient, necessary, sufficient, default,* and *evidential*. In this paper, we focus on the use of the latter three types. *Necessary* can be represented as a *sufficient* warrant with antecedent and conclusion reversed, while *necessary and sufficient* is equivalent to two *sufficient* warrants, one in each direction, between antecedent and conclusion.

Finally, we discuss how the type of warrant and the type of reasoning step are combined to yield the qualification for warrant application aspect of the tau backing. We summarize the interaction between warrant type and reasoning step as follows:

warrant type		**reasoning step**	**tau backing**
sufficient		MP, MT	sure
sufficient		ABD, ABC	possible
	-->$_s$		
default		MP, MT	usual
default		ABD, ABC	possible
	-->$_d$		
evidential		MP, MT	probable
evidential		ABD, ABC	possible
	-->$_e$		

As in deductive logic, a *modus ponens* or *modus tollens* reasoning step with a *sufficient* warrant yields a *sure* qualification on the warrant application. A *default* warrant differs from an *evidential* warrant in that the default warrant assumes a higher, i.e., *usual* -vs.- *probable*, qualification on warrant application using *modus ponens*. Reasoning abductively, essentially from conclusion to antecedent, will, in general, lead to less strong qualifications on the warrant application. For example, a warrant that is *sufficient* in one direction is considered less than *evidential* in the other direction. Reasoning with *any type* warrant abductively leads only to *possible* conclusions.

IV. Generating Argument Structures

We next describe a recursive algorithm for generating argument structures, given a claim, a set of domain knowledge (as warrants and data), and a set of inputs (as data). Claims, for now, may not be warrants. In Toulmin's terminology, the arguments generated are warrant using rather than warrant-establishing arguments. For any given claim, an argument is generated in a backward-directed fashion by looking for backing for the claim. A claim has atomic backing if it is already present in the given knowledge or data bases. Tau backing is generated by searching for warrants that are relevant to the claim, i.e., containing the claim or its negation in the antecedent or conclusion of the warrant. For each relevant warrant, a new tau structure is generated; new claims, from the other "side" of the warrant, are added to a global claims list. Any loop (i.e., a claim being used to support itself in an argument) or contradiction (i.e., the negation of a claim being used to support the claim) is pruned during the argument generation process.

The argument algorithm then calls itself recursively, attempting to find backing for the next claim on the claims list. The process continues until the claims list is empty. The algorithm then restarts the argumentation process for the claim that is the negation of the original claim. Final output is the two argument structures, i.e., claims and taus for both the input claim and its negation.

This algorithm has been implemented in a Scheme program. It has been tested on a number of small examples drawn from the literature in both argumentation and AI. We present program input and graphical output generated by the program for several of these examples in the next section.

V. Examples

We give several examples of argument structure generation that highlight different facets of the theory. The examples include arguments having knowledge that is incomplete, uncertain, or inconsistent, using different inference types, and involving default knowledge. We display several argument structures generated by our current implementation of the theory, highlighting relevant aspects of the arguments presented.

<u>Example 1 - inconclusive arguments</u>

Here we revisit the "Bermuda" argument of the introduction, represented as warrants and input data as follows:

$$(w1\ (\ (british\ passport)\ )\ \dashrightarrow_e\ (\ (british\ citizen)\ )\ sure\ given\ )$$

(w1 ((british passport)) -->$_e$ ((british citizen)) sure given)
(w2 ((bermuda born)) -->$_d$ ((british citizen)) sure given)
(w3 ((english speaking) (bermuda SSN)) --> e ((bermuda born)) sure given)
(w4 ((bermuda ssn)) -->$_s$ ((working papers)) sure given)
(w5 ((bermuda born) (alien parents)) -->$_s$ ((not british citizen)) sure given)

(d1 (english speaking) sure given)
(d2 (not british passport) sure given)
(d3 (working papers) sure given)

We generate argument structures that summarize the evidence for and against the possibility that Harry is a British citizen, as follows:

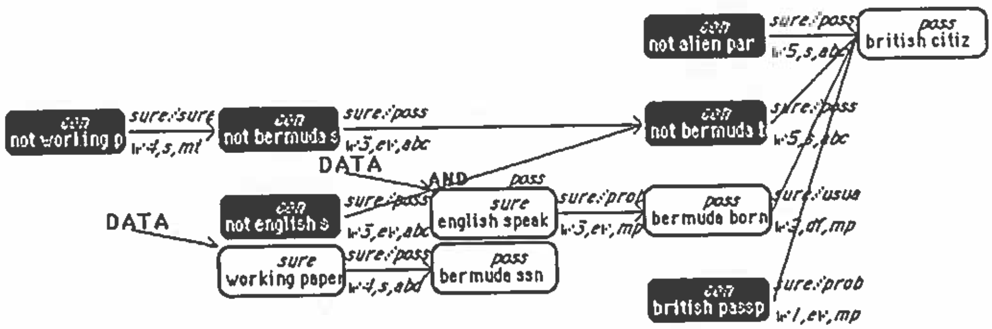

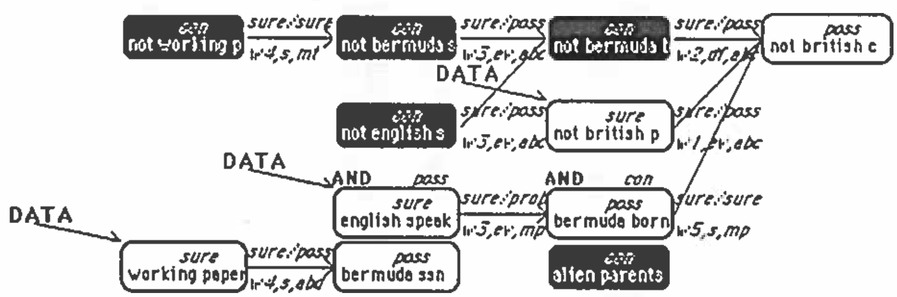

Claims are shown as nodes and warrants are represented as arcs between them. The notation under the arcs indicates the warrant, warrant type, and reasoning step type; notation above the arcs gives the warrant and warrant application qualifications. Notation in the claim nodes gives the statement of the claim and its qualification. Those claim nodes that are contingent, i.e., without supporting input data, are darkened.

In this example, we see that there is weak support for each of two mutually exclusive claims, british citizen and not british citizen. The strongest argument for british citizenship, that Harry was born in Bermuda, has a weak link at "bermuda ssn". From warrant W4 we see that "working papers" could be concluded from "ssn", but here the warrant is used abductively and the qualification reflects uncertainty resulting from this form of reasoning. That is, there may be other ways of obtaining working papers than having a ssn. Since the knowledge base is silent on these other ways - if any - belief in ssn may be strengthened. However, it remains that abductive reasoning is less certain than deductive reasoning.

In the not british citizen argument tree, the negative claim is supported by the data that Harry does not have a british passport. Since this is based on an evidential warrant, and the reasoning type is non-deductive, it is only weak support. Also in this argument tree the negative side has an opportunity to argue strongly against the claim, and simultaneously defeat the strongest argument for the claim, by showing that Harry's parents were aliens. Given the current knowledge base, this argument is only a hypothetical one.

Example 2 - inconsistent arguments

Next we consider the "republican-quaker-hawk-dove" example, similar to that discussed by Poole [19]. We present the warrants and input data, followed by the argument structures generated for dove and ~dove, as follows:

(w1 ((republican)) -->$_e$ ((not dove)) sure given)

(w2 ((quaker)) -->$_d$ ((dove)) sure given)

(w3 ((not dove) -->$_e$ ((supports star wars)) sure given)

(w4 ((quaker)) -->$_s$ ((religious)) sure given)

(d1 (religious) sure given)

(d2 (republican) sure given)

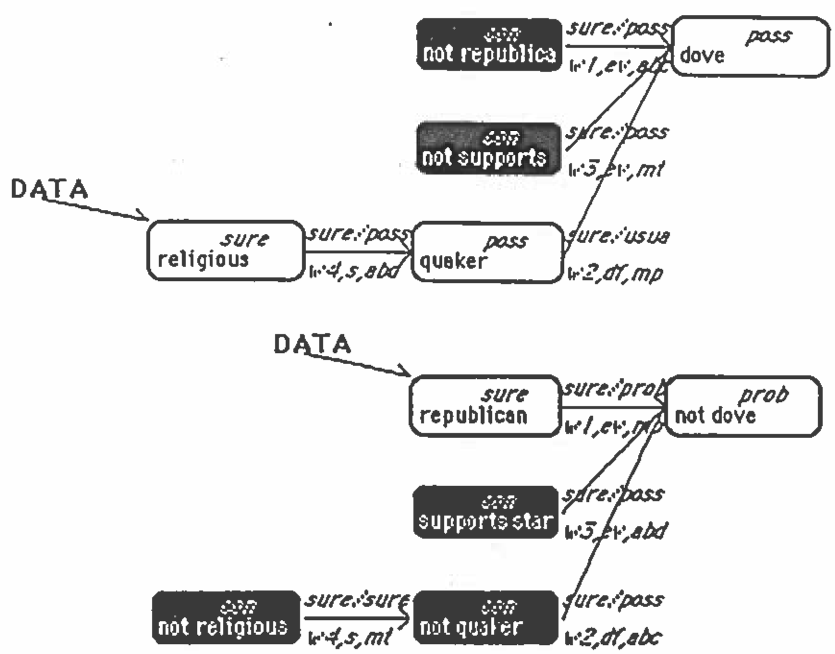

Here we see that our qualification on "~dove" is stronger, as *probable* due to the evidential support from "republican", while the support of religious only supports a belief in "quaker" as *possible*, which then can only support "dove" as *possible*. If we knew that the person was a "quaker" rather than just "religious", as the problem is usually presented, then the qualification for "dove" would be stronger, as *usual* due to the default warrant w2. If w2 is considered only evidential, then both claims would be established with a qualification of *probable* - "dove" because of the evidential support from "quaker", and "~dove" because of the evidential support from "republican".

Both of the last two arguments above are inconsistent in that there is no consistent assignment of probabilities to the two outcomes, as satisfying the qualitative qualifications would imply both have probability above 0.50 [8].

Knowledge about star-wars support would bolster one or the other of these two arguments.

Example 3 - default reasoning

Our version of the standard penguins-birds-flying example, illustrating how argumentation would handle standard default reasoning issues as encountered by inheritance-based representations, is as follows:

(w1 ((bird)) -->$_{d}$ ((flies)) sure given)

(w2 ((penguin)) -->$_{s}$ ((not flies)) sure given)

(w3 ((penguin)) -->$_{s}$ ((bird)) sure given)

(d1 (bird) sure given)

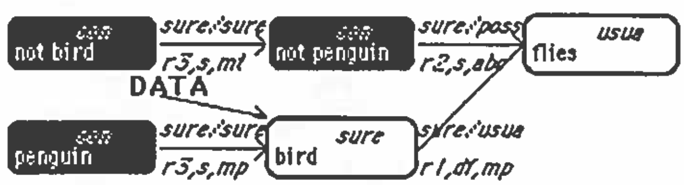

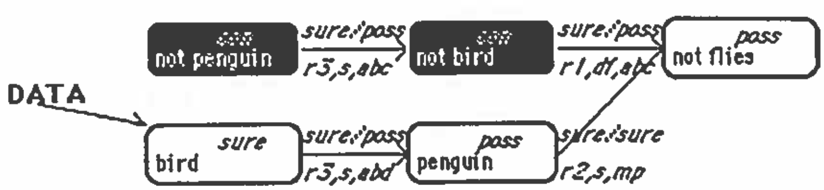

Given the data as "bird", we find that the strongest argument would be for "flies", with qualification *usual* due to the default warrant w1, while "~flies" would only be deemed *possible*, due to the weak abductive argument that supports penguin. On the other hand, if we had the data "penguin", the claim "~flies" would be established with qualification *sure* due to the sufficient warrant w2, while "flies" would be established with a qualification of *usual*. While the arguments are inconsistent, as discussed above, there is a preference for one conclusion (the normally preferred one of "~flies") over the other.

Here our broad quantification of belief levels is sufficient to draw the expected results (given the stronger claim is taken to be the result) for this kind of implicit default reasoning. However, if warrant w2 were only taken as default information about penguins, we would not be able to distinguish between the two argument strengths. Again, we see the impact of warrant type decisions on argument outcomes.

Example 4 - multiple causes

Here we look at a multiple causes situation involving rain, sprinkler, and wet grass [16]. We generate some troublesome results, which motivates the need to address argumentation as an interaction between arguments, as we plan to explore in the next phase of our research. The specification is as follows:

(w1 ((rain)) $-->_s$ ((wet grass)) sure given)

(w2 ((sprinkler)) $-->_s$ ((wet grass)) sure given)

(w3 ((wet grass)) $-->_s$ ((wet shoes)) sure given)

(d1 (sprinkler) sure given)

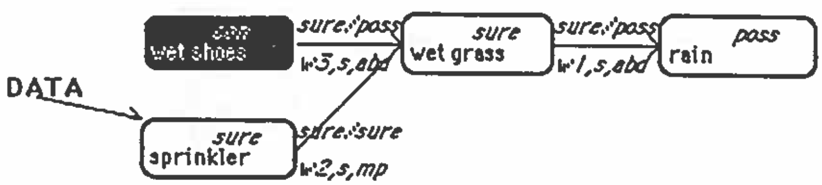

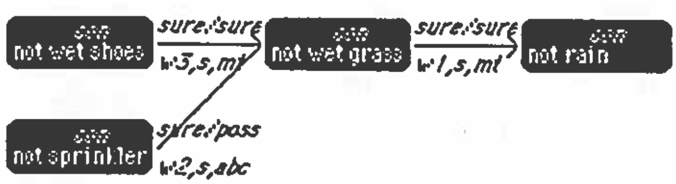

Here, the warrants are representing causal relationships. Diagnostic reasoning often reflects weak abductive reasoning over such "causal warrants". Here "rain" is defeasibly supported by "wet grass", which is strongly supported by the "sprinkler" input data. This last part is odd in that "sprinkler", known to be true, provides a causal explanation of "wet grass" and would undercut "rain" rather than support it, as the argument generated here does.

VI. Related Research

Our research in modeling argumentation builds on other work in argumentation from both inside and outside AI. Also, since we have chosen to explore argumentation as a method for reasoning in weak theory domains, this work overlaps somewhat with work in AI on reasoning under uncertainty. We discuss these related areas in this section.

There has been a great deal of investigation of argument outside AI, for example, in philosophy, rhetoric, and critical reasoning (e.g., [4], [7], [10], [11], [22], [26]). Our view of argument as a combination of explanation, or "proof", and dialectic is based on two senses of argument found in this literature (e.g., [4]). Our representation of argument structure is based on Toulmin's well known data-warrant-claim model of argument [26]. The idea that argumentation is an important reasoning technique is discussed in [7], [11], [22], and others. It is mainly our formal approach to modeling argument that differentiates us from this other work. At least one investigator [9] has called for more formality in argumentation research: "[modeling arguments] will be beneficial in several respects...models yield a precise and clear understanding or the exact character of the theories they serve". One of the main goals of the current work is to employ the computer as a modeling tool to achieve just this effect. The dynamic nature of computer programs supports research in argumentation generation as well as argument analysis.

AI researchers have also undertaken work in argumentation, and recently research in this area seems to be on the rise [21]. Most AI work in argumentation has been done in contexts other than reasoning per se, e.g., natural language processing (e.g., [6]), software engineering (e.g., [3], [13]), or legal reasoning (e.g., [1], [15], [24]). This previous work can be divided into work that is mainly concerned with argument analysis and structure ([3], [13], [15], [25]) and work that is mainly concerned with argument generation ([1], [6], [23]). Results from other work in AI and argumentation are important to our theory. In particular Marshall's and Storrs' work in the structure or argument is relevant, as is the work of Rissland, Ashley, and Flowers, et. al., in generating arguments and argument moves.

Our current research continues this work and adds to it, by (1) its emphasis on both the structure and generation of arguments; (2) use of a domain independent theory of argument; (3) incorporation of two senses of argument, supporting explanation and dialectical process; (4) implementation of the theory in a computer program, and evaluation of the results; and (5) analysis and representation of various types of uncertainty in the domain knowledge.

The last point brings this research into the realm of AI work in uncertain reasoning. Since we are investigating argumentation as a general approach to uncertain reasoning, it is worthwhile to briefly mention similarities and differences between this approach and other, better established approaches to uncertain reasoning. (Helpful surveys of these approaches can be found in [5] and [12].)

Argumentation is a general approach to uncertain reasoning that expects inconsistency, as well as incompleteness and uncertainty, in the knowledge, and handles it by generating arguments both for and against a claim. An argumentation approach emphasizes justifying decisions, rather than decision making per se. Justification for a claim includes both providing support for a claim and defending it with respect to support for other plausible claims.

In contrast, other work in uncertain reasoning has concentrated on uncertainty per se, rather than on incompleteness or inconsistency. Indeed, none of the methods mentioned can sensibly handle total conflict in the knowledge [5]. Also, the emphasis is on decision making, e.g., by choosing the claim with the highest certainty, rather than on decision justification. If justification is included, it consists of knowledge that supports the claim, but ignores weak points in the support and support for alternative claims (e.g., [20]).

In our theory, argumentation is independent of the various approaches to reasoning under uncertainty. Some method for representing, combining, and propagating uncertainty is presumed, but it can be numeric or symbolic, probabilistic or not, and so on. For the purposes of implementing our theory, we invented a simple, symbolic representation of uncertainty and an algorithm for propagating it throughout an argument. This method captures uncertainty introduced by different reasoning types, as well as uncertainty in the knowledge itself. Argumentation could make use of any method of reasoning under uncertainty.

In sum, the current work, along with other AI work in reasoning under uncertainty, assumes the existence of uncertainty in the domain knowledge. Uncertainty is broadly defined to include incomplete and inconsistent as well as uncertain knowledge. Argumentation, which emphasizes decision support and justification, is seen as a useful, general technique for reasoning in light of the problems that arise from reasoning in weak theory domains.

VII. Conclusion

We have presented a formal theory of argumentation as a method for providing decision support and justification for plausible reasoning in weak theory domains. Two senses of argument, argument as supporting explanation and argument as interactive process, are identified. The complete theory is to include a structure for representing and an algorithm for generating each of these types of argument. We described a partial implementation of the theory, a program, ART, that generates supporting arguments for a given claim and counter-claim. ART uses plausible reasoning (e.g., abductive reasoning) in addition to deductive reasoning (i.e., modus ponens and modus tollens). Hypothetical reasoning is also incorporated to point out incompleteness in the knowledge. A simple method for representing, combining, and propagating the uncertainty introduced by plausible reasoning is given; the same method is used to represent uncertainty in the knowledge base. The algorithm itself can be combined with any general method for handling uncertainty. Inconsistency in the knowledge base is expected, so ART always attempts to generate support for both a claim and its negation. As yet, there is no way to directly compare support for alternative claims. We gave examples of program output for a set of paradigmatic problems from weak theory domains.

Next, we plan to design and implement a formal theory of argument as an interactive process. The next stage will use the format of a two-sided argument to intertwine the strengths and weaknesses of support for competing claims so they can be directly compared to one another.

For example, as noted in example 4, the pro side could present its argument for rain, that the grass is wet and rain is a cause of wet grass. But the con side can undercut this argument by pointing to an alternative explanation for the wet grass, namely, that sprinklers also cause wet grass, and the sprinkler was indeed on. It is this type of control, selecting particular argument moves [23], that we will pursue in the next phase of our research.

In example 3, though flies is highly likely based on the given data bird, an arguer would look for the weakness in the argument, and seek verification that the bird is not a penguin. Hypothetical lines of reasoning will also be explored, to point out additional knowledge that, if available, could bolster a claim or resolve an apparent inconsistency.

Hypothetical arguments can be based upon the contingent (darkened, in our figures) aspects of arguments. These contingent elements could serve as the basis for the generation of information gathering goals, if the argument process were embedded in an entity that solves problems in weak theory domains.

Issues that must be addressed in future work include (i) identifying a structure for interactive argument; (ii) identifying a process for generating this structure that permits intertwining the tasks of supporting and defending a claim, along with rebutting and undercutting competing claims; (iii) identifying argument moves and heuristic strategies for controlling the argument generation process.

References

[1] Ashley, K., Toward a computational theory of arguing with precedents, in: *Proceedings ICAIL-89*, Vancouver, Canada (1989) 93-102.

[2] Buchanan, B. and E. Shortliffe, *Rule-Based Expert Systems* (Addison-Wesley, 1984).

[3] Conklin, J., Design rationale and maintainability, MCC TR-STP-249-88, 1988.

[4] Cox, J. R. and C. A. Willard, eds., *Advances in Argumentation Theory and Research* (Southern Illinois University Press, 1982).

[5] Dubois, D. and H. Prade, Handling uncertainty in expert systems - pitfalls, difficulties, remedies, in: Hollnagel, E., ed., *The Reliability of Expert Systems* (Ellis Horwood Limited, 1989) 64-118.

[6] Flowers, M. and R. McGuire & L. Birnbaum, Adversary arguments and the logic of personal attacks, in Lehnert, W. and M. Ringle, eds., *Strategies for Natural Language Processing* (Erlbaum Associates, 1982).

[7] Freeley, A., *Argumentation and Debate: Critical Thinking for Reasoned Decision Making*, 7th ed., (Wadsworth Publishing Company, 1990).

[8] Goldszmidt, M. and J. Pearl, On the consistency of defeasible databases, *Artificial Intelligence* 52 (1991) 121-149.

[9] Hample, D., Modeling argument, in Cox, J. R. and C. A. Willard, eds., *Advances in Argumentation Theory and Research* (Southern Illinois University Press, 1982).

[10] Horner, W., *Rhetoric in the Classical Tradition* (St. Martin's, 1988).

[11] Kahane, H., Logic and Contemporary Rhetoric, (Wadsworth Publishing Co., 5th Ed., 1988)

[12] Lea Sombe, Reasoning under incomplete information in artificial intelligence, *International Journal of Intelligent Systems* 5 (1990) 323-472.

[13] Lee, J., DRL: A task-specific argumentation language, in: *Proceedings AAAI Spring Symposium on Argumentation and Belief*, Stanford, CA (1991) 122-132.

[14] Levi, E., *An Introduction to Legal Reasoning* (University of Chicago Press, 1949).

[15] Marshall, C., Representing the structure of a legal argument, in: *Proceedings ICAIL-89*, Vancouver, Canada (1989), 121-127.

[16] Pearl, J., Embracing causality in formal reasoning, in: *Proceedings AAAI-87*, Seattle, WA (1987), 369-373.

[17] Pollock, J., Defeasible reasoning, *Cognitive Science* 11 (1987) 481-518.

[18] Polya, G., *Patterns of Plausible Inference* (Princeton University Press, 2 ed., 1968).

[19] Poole, D., Explanation and prediction: an architecture for default and abductive reasoning, *Computational Intelligence* 5 (1989) 97-110.

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

[21] *Proceedings AAAI Spring Symposium on Argumentation and Belief*, Stanford, CA (1991).

[22] Rescher, N., *Dialectics: A Controversy-Oriented Approach to the Theory of Knowledge* (State University of New York Press, 1977).

[23] Rissland, E., Argument moves and hypotheticals, in: Walter, C., ed., *Computing Power and Legal Reasoning* (West Publishing Co., 1985).

[24] Rissland, E., Dimension-based analysis of hypotheticals from Supreme Court oral argument, in: *Proceedings ICAIL-89*, Vancouver, Canada (1989) 111-120.

[25] Storrs, G., Extensions to Toulmin for capturing real arguments, in: *Proceedings AAAI Spring Symposium on Argumentation and Belief*, Stanford, CA (1991) 195-204.

[26] Toulmin, S., *The Uses of Argument* (Cambridge Univ. Press, 1958).