Burden of Proof in a Computational Model of Argumentation

Arthur M. Farley
Kathleen Freeman

CIS-TR-94-17
July 1994

Department of Computer and Information Science
University of Oregon

Burden of Proof in a Computational Model of Argumentation

Arthur M. Farley
Kathleen Freeman

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

ABSTRACT

In this paper, we present a computational model of dialectical argumentation. As a structure of supports and rebuttals, an argument consists of a series of qualified claims and their backings. When viewed as a process of interacting moves, an argument consists of move generation functions and selection heuristics. We introduce a concept of burden of proof which acts as filter on applicable moves and as termination criterion, determining when an argument has been completed and which side has won. Burden of proof is an important component of argumentation when used as a means for making decisions in the context of uncertain, incomplete, and conflicting knowledge.

INTRODUCTION

We present results of our research identifying and implementing a computational model of dialectical argumentation that exhibits reasonable behavior when applied to benchmark examples from formal argumentation and artificial intelligence (AI) research. The work was motivated by the desire and need to incorporate the ability to construct arguments in programs that attempt to model intelligent behavior and to explore the use of dialectical argument as basis for making decisions in domains where knowledge is incomplete, uncertain, and inconsistent, i.e., *weak theory domains.*

In a weak theory domain, a reasoner cannot count on deriving claims that are deductively valid. Since claims cannot always be proved conclusively, a definition of "proof" is needed that makes sense for weak theory domains. Argumentation, with its emphasis on both support and refutation of claims allows for the definition of new concepts of proof. For example, we will say a claim is *dialectically valid* (with respect to a given set of knowledge) when there is plausible, irrefutable support for a claim and no such support for a counter-claim. We define additional "proof levels" by constraining or weakening the conditions of dialectical validity. As an example of a weaker proof level, to prove a claim with a *preponderance of the evidence*, there may be support for both a claim and a counter-claim, as longer as there is more or stronger support for the input claim.

Explicit representation of proof levels is relevant because it corresponds to our experience with real world problems, where not all claims need be proved to the same level to warrant action. Proof levels are an important element of *burden of proof*, defined as which side to an argument must realize what proof level. One area where burden of proof is used is the legal domain. Different burdens of proof are required at different stages of the legal process and for different types of legal action. We will see that burden of proof is a useful component of a computational model of argumentation. It will serve to focus the argument and improve the efficiency of argument generation.

The concept of burden of proof as introduced above has been incorporated into a computational model of argument. 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. In the following, we first describe our model of argument then discuss the concept of burden of proof with respect to the model. We demonstrate the application of different burdens of proof in a simple argument, providing trace fragments of output from the computer program that implements the model and graphical representations of argument structures.

MODELING DIALECTICAL ARGUMENTATION

The representation of arguments as structured entities and the generation of arguments as dialectical processes are both crucial to our theory. For argument as supporting explanation, we create argument structures that serve to organize relevant, available, plausible support for a claim and for its negation. Argument as dialectical process includes the tasks of supporting and refuting claims, and choosing actions relevant to these tasks.

We represent a supporting argument in an extended version of the form given by Toulmin in *The Uses of Argument* (1958). For Toulmin, an argument comprises *data* (i.e., evidence, grounds) said to support a *claim* (i.e., conclusion). The authorization for moving 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 conclusively, i.e., the resultant claim may be *qualified*. The claim may be subject to *rebuttals*, special circumstances where the warrant would not hold.

In our representation of argument, the data and warrant are also seen as claims, and therefore subject to rebuttal and qualification. (Since all the major elements of an argument are claims, we will refer to these as data, warrant, and *conclusion*, to avoid ambiguity.) All claims (not just warrants) must be supported, i.e., have backing. We define two types of backing: *atomic*, as information from outside the domain of argument, and *tau* ("Toulmin argument unit"), where the conclusion is supported by a data and a warrant. A claim may have multiple backings. A rebuttal is a rival claim (currently defined as the negation of the claim) and the arguments that support the rival conjecture.

Since warrants represent a relationship between two claims, they have a structure that differs slightly from other claims. In addition to qualification, backing, and rebuttal, a warrant has two propositional fields, *antecedent* and a *consequent*. The antecedent and consequent fields consist of one or more propositional clauses. Multiple clauses in either the antecedent or consequent are assumed to represent conjunctive clauses.

A warrant also has two type fields. The *wtype1* field classifies the relationship between the antecedent and consequent as *explanatory (ex)* or *sign (si)*. This distinction is made in (Freeley, 1990), for example. An example of an explanatory relationship is a causal link, because knowledge of the antecedent "explains" knowledge of the consequent, e.g., antecedent "fire" causes (explains) its consequent "smoke" (see, e.g., Porter, et. al., 1990). A sign relationship represents a correlational link between antecedent and consequent, for example, "Most plane crashes occur on Tuesdays." This type distinction is important in determining possible moves during dialectical argument generation.

The *wtype2* field of a warrant represents the strength with which its consequent can be drawn from the given antecedent. Representing information as to the strength of the connection between warrant fields is appropriate for reasoning with incomplete or uncertain knowledge. Current types are *sufficient (s)*, *default (df)*, and *evidential (ev)*. The sufficient type is meant to represent certain relationships, e.g., definitions. Default and evidential are meant to represent two levels of uncertain knowledge, with default indicating relationships that are usually the case (e.g., "birds fly"), and evidential referring to less certain links (e.g., "persons who live in Bermuda are often British subjects"). Warrants are expected to be written in the direction that accommodates the strongest possible type.

TABLE 1. Reasoning Types

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)

An important aspect of our model is the application of warrants in various "directions". Given a warrant with antecedent p and consequent q, we define allowable *reasoning steps* in Table 1. The last two reasoning steps are fallacies in the context of deductive reasoning (asserting the consequent and denying the antecedent, respectively). However they are appropriate for reasoning when knowledge is incomplete or uncertain. Polya (1968) discusses similar "patterns of plausible inference". They are "examining a ground" (MP, ABC in Table 1) and "examining a consequent" (MT, ABD in Table 1).

When deductive and plausible reasoning types are present in the same system, as they are in this model, 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 abductive reasoning would allow the reasoner to derive the conclusion "sprinkler on" from the data "rain". To permit the generation of acceptable conclusions, while blocking generation of unacceptable ones, the reasoning types interact with the warrant type fields. Modus ponens/abduction combinations are not permitted for two explanatory warrants, unless both warrants are "evidential".

Qualifications are needed to capture the support for claims reached as a result of arguments using uncertain knowledge and plausible reasoning. We use the following qualifications: *strong(s)*, *usual (!-)*, *credible (+)*, and *unknown (?)*. The first three are ranked in decreasing order of support; the last indicates a lack of (known) support.[1]

The qualification on a claim is that associated with its strongest supporting argument. The qualifications on input data and warrants are given as atomic backing at input time and remain unchanged thereafter, unless better support is derived from a tau backing. The qualification on

[1]The issue of representing, combining, and propagating uncertainty is a research area in it own right. Here, we choose a simple representation that covers our needs with respect to the argumentation model, but avoids the complexity and commitments that would come with implementing one or another of the current methods for representing uncertainty.

TABLE 2. Link Qualifications

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

any claim resulting from a tau backing is the least of the qualifications associated with the warrant application: the qualification(s) on the data support, the qualification on the warrant, and the qualification derived from the warrant type and reasoning step applied ("link qualification", see Table 2). This weak link approach to propagating support for claims and its appropriateness for plausible reasoning is discussed in (Pollock, 1991) and (Rescher, 1976).

We have just described a representation for arguments as supporting explanations. Having only a structural model does not capture the procedural character of argumentation. Generating dialectical arguments results in the intertwining of the argument structures that support a claim and its negation. Support for the input claim is *Side-1 support* and support for the negation of the claim is *Side-2 support*. Support for a claim is *consistent*, for mutually exclusive conjectures, if there is support on one side of the claim and no support on the other side, i.e., the qualification slot for one side only contains a "!", "!-", or "+". A claim that is not consistent is *controversial*. A claim qualification is *terminal* when there are no argument moves available that could change its qualification.

Dialectical argument begins with Side-1 attempting to find support for its claim. If there is no support, the argument ends, and Side-1 concedes the claim. If Side-1 is able to find support for the claim, control is given to Side-2, which tries to refute the argument or claim(s) established by Side-1. A side executes argument moves until one succeeds or there are no more moves. When a move is successful, the qualification on the input claim becomes a *check* (i.e., as in chess) for the current side. The current side gives control of the argument to the other side, to attempt to refute the refutations. When there are no more moves, the current side concedes the argument.

Finding support for a claim results in the generation of the argument structures described earlier in this section. Given a claim, search for support proceeds in a goal-directed fashion by looking for backing for the claim. The process grounds out when a claim is supported in the data base. Then a new tau structure is generated for each warrant that supports a claim, and the qualification and backing fields of the claims are updated to reflect the new support.

Next, we focus our attention on refuting a claim or its supporting arguments. We begin to give operational meaning to the tasks of dialectical argumentation by defining them in terms of the argument moves needed to implement them (see Table 3). We distinguish two types of refutation: (a) *undercutting* and (b) *rebutting* (following Pollock, 1987, though not exactly). Undercutting is accomplished by finding weaknesses in purported support for a claim. With respect to the structure of a tau, undercutting questions the sufficiency of the data support and the link fields (i.e., warrant type and reasoning type).

TABLE 3. Dialectical Argument Moves

ARG TASKS	MOVES	GIVEN	SHOW	SUCCESS C IS -
support C	(a) support	C	X->C ^ X C->X ^ X ~C->X ^ ~X X ->~C ^ ~X	supported supported supported supported
refute C undercut C	(b) invalid data	X->C ^X	~X	controversial
	(c) exception	X->$_{df}$C ^X	X^Y->~C ^Y	not supported
	(d) inapplicable evidence	X->~C ^~X	Y->s~C ^Y	not supported
	(e) unneeded explanation	C->exX ^X	Y->exX ^Y	not supported
	(f) missing evidence	X->$_{ev}$C ^X	Y->$_{ev}$C ^~Y	controversial
	(g) conflicting evidence	X->$_{ev}$C ^X	Y->$_{ev}$~C ^Y	controversial
rebut C	(h) reductio ad absurdum	C	C->...->Z ^~Z	controversial
	(i) support rival claim	C	X->~C ^X	controversial

If an undercutting move is successful, it can result in withdrawal of an argument. Also, in Table 3, we see that some methods for which undercutting a claim may result in support for a rival claim. These leave the original argument as it was, but highlight its inadequacies. As a result, its conclusion may become controversial or unsupported.

In contrast, rebutting moves attack a claim without regard for its supporting arguments. In a successful rebuttal, support for a rival claim is found. As a result, the original claim becomes controversial.

When a side is in control of the argument, it must select which argument move to apply. Heuristics, guidelines for ordering argument moves, determine the course of the actual argument. They are meant to reflect two goals: generate the strongest argument possible (for each side), and generate a coherent argument, i.e., the arguments of each side are responsive to the arguments put forward by the other side.

Heuristics are used to order both the moves that implement a dialectical argumentation task, and the warrants that implement a particular move. Argument moves are ordered as follows: (a) moves that can defeat a tau are preferred over moves that can only cause the tau's claim to become controversial; (b) specific moves are preferred over general ones; (c) moves that attack taus closer to the root are preferred; and (d) undercutting moves are given priority over rebutting moves.

Warrants are currently ordered according to the following criteria: (a) strong reasoning types are preferred; (b) strong warrant types are preferred; and (c) warrants where the data support part already has consistent support are preferred. These heuristics anticipate to some extent the moves that the other side may use to try to refute this tau, should this warrant turn out to actually support the claim. That is, strong reasoning types cannot be refuted, while weaker ones can. Warrant type *sure* cannot be refuted; warrant type *default* can only be refuted by an exception move; while evidential warrants are refuted by both conflicting evidence and missing evidence moves. Similarly, weaker reasoning types allow more opportunities for rebuttals. Controversial or negated data can support a claim controversially at best, while consistently supported data or data where nothing is known about the support will be, at least immediately, less susceptible to the question data support argument move.

This completes our overview of our computational model of argumentation. Given a set of warrants, input data, and claim, our implemented system proceeds to generate a dialectical argument. Control switches from side to side as check qualifications are realized by successful moves. Deciding which moves, if successful, are sufficient to generate a check qualification for a side and deciding when the argument process is complete depends upon the burden of proof.

BURDEN OF PROOF

There is one element missing in our model of argumentation as described above: the notion of burden of proof, its implementation, and its impact on arguments generated. There are two aspects to the notion of burden of proof: (1) which side bears the burden; (2) what is the level of proof required. As for which side bears the burden, as we consider only two sides, this must be either *Side-1* or *Side-2*. In our discussion to follow, we assume that *Side-1* always bears the burden of proof and that the claim is stated in terms of that side's perspective.

We now define a set of proof levels in terms of our model. A *dialectically valid (dv)* claim is one with terminal, consistent support. This implies all competing arguments have been defeated, leaving a surviving explanation for the goal claim. When a claim has terminal support, either consistent or controversial, that is stronger than the support for its negation (if any), the claim is proved with a *preponderance of the evidence (pe)*. When there is terminal support for a claim, either consistent or controversial, regardless of the strength of the support, the claim is proved with a *scintilla of evidence (se)*. Finally, the burden of proof level *beyond a reasonable doubt (brd)* is realized by support for a claim that is both dialectically valid and strong (i.e., qualified as sure or usual). As noted, we consider the burden of proof in an argument to be on Side-1, i.e., to win an argument Side-1 must establish adequate support for the input claim, where adequate support is defined with respect to one of the proof levels given above.

Burden of proof plays several roles in the process of argumentation:

- as basis for deciding the relevance of particular argument moves;
- as basis for declaring an argument to be finished;
- as basis for deciding the outcome -- winner -- of an argument.

We will demonstrate these roles and illustrate our model of argumentation by considering a series of examples based upon a classic problem in the argumentation literature (Toulmin 1958):

> Anyone born in Bermuda can be assumed to be a British subject, unless both parents were aliens. Statistics show that most people who speak English and have Bermuda identification numbers were born in Bermuda. Having a Bermuda identification number makes one eligible to obtain Bermuda working papers, as does having special working skills, unless there are already enough workers with the special skill. Most people who have a British passport, or pay taxes in Bermuda, or live in Bermuda, are British subjects.

This knowledge can be represented by the following warrants:

```
(w1 ((british passport)) --> ex ev ((british subject))
(w2 ((bermuda born)) --> ex df ((british subject))
(w3 ((english speaking)(bermuda id#)) --> si ev ((bermuda born))
(w4 ((bermuda id#)) --> ex df ((working papers))
(w5 ((bermuda born)(alien parents)) --> ex df ((not british subject))
(w6 ((special skills)) --> ex df ((working papers))
(w7 ((special skills)(quota met)) --> ex df ((not working papers))
(w8 ((pays Bermuda taxes)) --> ex ev ((british subject))
(w9 ( (lives in Bermuda)) --> ex ev ((british subject))
```

For this particular problem, we are given the following input data regarding Harry:

> Harry is English speaking, lives in Bermuda, and pays Bermuda taxes. He does not have a British passport, but has working papers. He has special work skills, though the quota for working papers for Harry's special skills has been met. We would like to know whether or not Harry is a British subject.

This situation can be represented by the following input claims:

```
(d1 (english speaking) (!? OBSERVED))
(d2 (british passport) (?! GIVEN))
(d3 (special skills) (!? GIVEN))
(d4 (quota met) (!? GIVEN))
(d5 (pays Bermuda taxes) (!? GIVEN))
(d6 (lives in Bermuda) (!? GIVEN))
(d7 (working papers) (!? GIVEN))
```

We would like to know whether or not Harry is a British subject, represented as:

```
(claim (british subject) (?? NIL))
```

In the first example, the burden of proof to be borne by Side-1 will be scintilla of evidence. The argument starts with Side-1 attempting to generate support for the input claim. As the claim is not directly supported by inputs, the warrant base is searched for warrants that can be used to provide tau support for the claim. The resultant set of warrants is then heuristically ordered according to the criteria given above.

Here, warrant W2 is preferred over the other warrants because it is a default, rather than evidential, warrant. Warrants W8 and W9 are preferred over W1, because using W1 would mean using data support "British passport", and "not British passport" is directly supported in the current data base. The search proceeds as indicated in the program output below:

```
Looking for support for claim (BRITISH SUBJECT)
Check warrant(s) (W2 W8 W9 W1) for support for (BRITISH SUBJECT).

    Trying W2 with ((AND (BERMUDA BORN)))
    to support (BRITISH SUBJECT).
```

Looking for support for claim (BERMUDA BORN)
Check warrant(s) (W3) for support for (BERMUDA BORN).

Trying W3 with ((AND (ENGLISH SPEAKING) (BERMUDA ID\#)))
to support (BERMUDA BORN).

Looking for support for claim (ENGLISH SPEAKING)
Support found in data base for (ENGLISH SPEAKING)

Looking for support for claim (BERMUDA ID\#)
Check warrant(s) (W4) for support for (BERMUDA ID\#).

Trying W4 with ((AND (WORKING PAPERS)))
to support (BERMUDA ID\#).

Looking for support for claim (WORKING PAPERS)
Support found in data base for (WORKING PAPERS)

Support found for warrant W4
Support found for warrant W3
Support found for warrant W2

So, Side-1 finds support for the claim and generates an argument in the modified Toulmin format, which can be graphically represented as follows:

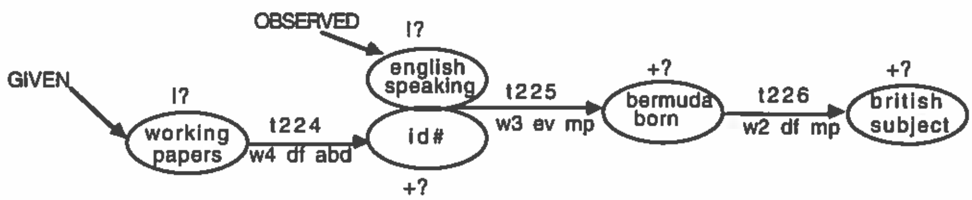

Though there is support for the top claim, it is only qualified as credible, rather than strong. The reason for this is that Side-1 has made use of evidential rules as well as abductive reasoning in generating the support for "British subject". Since Side-1 has achieved a Side-1 check qualification on the top claim, control of the argument is given to Side-2. Side-2 proceeds by generating possible moves:

Find exception for rule (AND (BERMUDA BORN)) -> (AND (BRITISH SUBJECT))
Refuted tau would be #:\t226

Question data: find support or bolster (AND (NOT BERMUDA BORN))
Refuted tau would be #:\t226

Look for evidence that supports (AND (NOT BERMUDA BORN))
to refute (AND (BERMUDA ID\#) (ENGLISH SPEAKING)) evidence for (AND (BERMUDA BORN))
Refuted tau would be #:\t225

Look for evidence, other than (AND (BERMUDA ID\#) (ENGLISH SPEAKING)), that would support (AND (BERMUDA BORN)) but is unavailable.
Refuted tau would be #:\t225

Question data: find support or bolster
(OR (NOT BERMUDA ID\#) (NOT ENGLISH SPEAKING))
Refuted tau would be #:\t225

Look for another explanation for (AND (WORKING PAPERS))
to show (AND (BERMUDA ID\#)) is an unneeded explanation.
Refuted tau would be #:\t224

Question data: find support or bolster (AND (NOT WORKING PAPERS))
Refuted tau would be #:\t224

Find support for (AND (NOT BRITISH SUBJECT))
Refuted tau would be TOP

Although there are many argument moves, most of them cannot defeat Side-1's support for the input claim (they would make the support controversial). But to win the argument when the proof level for Side-1 is scintilla of evidence, Side-2 must show that there is no support for the input claim. Therefore, only moves that could result in an outright defeat of Side-1's support for the claim are relevant, and the rest are discarded for the time being.

The only two remaining moves, in sorted order, are

Find exception for rule (AND (BERMUDA BORN)) -> (AND (BRITISH SUBJECT))
Refuted tau would be #:\t226

Look for another explanation for (AND (WORKING PAPERS))
to show (AND (BERMUDA ID\#)) is an unneeded explanation.
Refuted tau would be #:\t224

Side-2 attempts each of the remaining moves in turn. There is no exception found to the default, but the second move is successful. The resultant argument, graphically represented, now looks as follows:

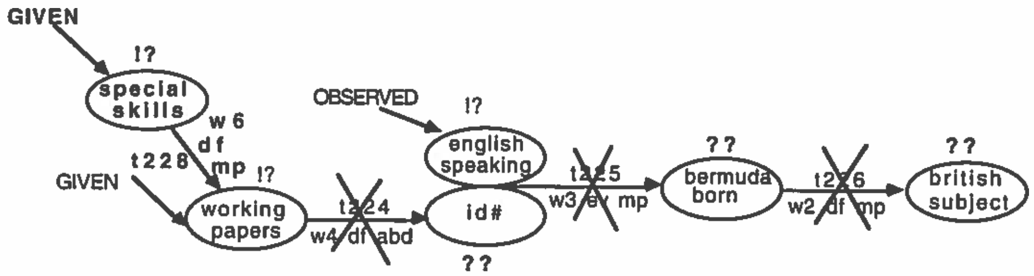

Side-2 succeeds in providing an alternative explanation for how Harry obtained his working papers, defeating the support for "id#". The support for "id#" is therefore withdrawn, as are the taus further up the argument tree that depended on "id#" for their support. Since its supporting argument was defeated, the qualification for "British subject" reverts to unknown. This is a check for Side-2 and control is returned to Side-1, which must attempt to defend the claim against this refutation or find another argument.

Since Side-2's new argument uses a default tau, Side-1 generates an exception move to defeat Side-2's defeater, along with a new support the top claim move. The exception move succeeds. A graphical representation of the resultant argument is as follows:

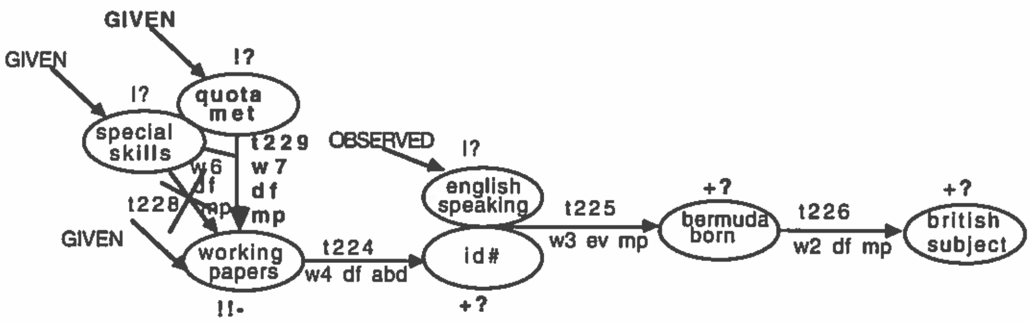

Side-1 succeeds on the first move: "special skills" does not provide an explanation for "working papers" after all, because the exception rule to the default holds in this case, i.e., the quota has been met. Side-1 has defeated Side-2's defeater argument, so Side-2's argument is removed, and the original argument is reinstated. The top claim is supported once again and control of the argument returns to Side-2. Side-2 is again faced with the task of defeating Side-1's argument directly, otherwise Side-1 has established its argument with a scintilla of evidence. Side-2's only move is the following:

Find exception for rule
(AND (QUOTA MET) (SPECIAL SKILLS)) -> (AND (NOT WORKING PAPERS))
Refuted tau would be #:\t229

There is nothing available in the data or warrant base to assist Side-2, and the move fails. Thus, Side-2 fails to defeat Side-1's argument. Side-1 has satisfied its burden of proof and has won this argument at the level of scintilla of evidence.

However, if we put the next higher burden of proof on Side-1, that of preponderance of the evidence, the argument would continue. To defeat Side-1 at this level, Side-2 need only show that there is at least as much evidence against the claim as there is for it. In addition to the (defeater) move which failed above, Side-2 would have all the controversy raising moves that were rejected earlier. Side-2 tries each move in turn. The only one that succeeds is finding an argument supporting the opposite of Side-1's initial claim. A graphical representation of the resultant argument is as follows:

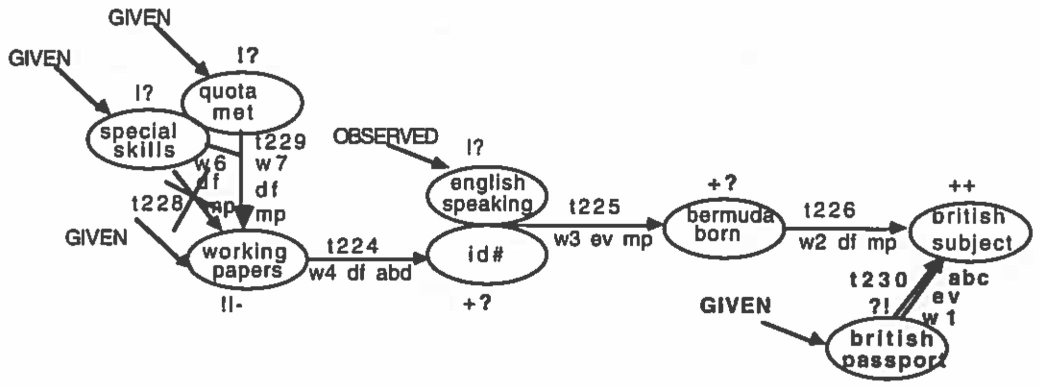

Side-2 does not succeed in refuting Side-1's last argument, but does manage to make the top claim controversial by bringing in new evidence, Harry's lack of a British passport. Since having a British passport is evidence for "British subject", this missing evidence tends (via the contrapositive of abductive reasoning) to support the counter-claim. The support is weak, but it is enough to create controversy over the top-level claim. Side-1 takes control of the argument. Side-1 can attempt to refute Side-2's new argument directly or find another line of support for its claim to outweigh Side-2's case. The latter is successful, as shown below:

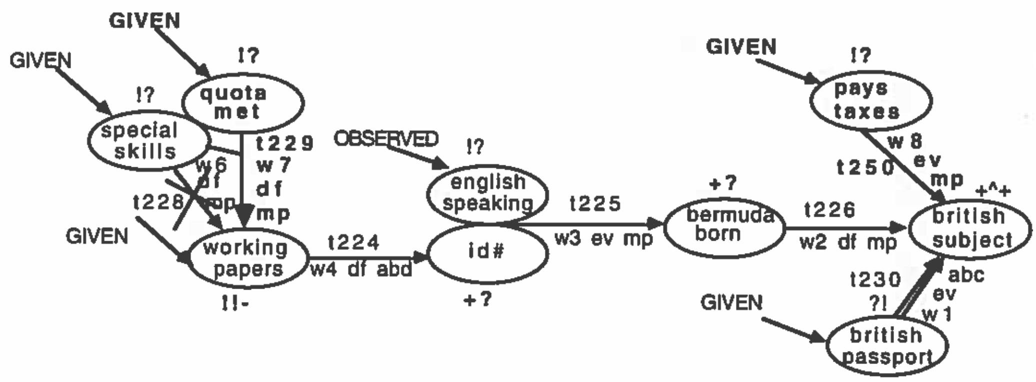

The qualification on "British subject" becomes (+^ +), where the annotation on the first field of the qualification means that the positive proposition has more evidential support than the negative proposition. Since the proof level is preponderance of the evidence, that is all Side-1 needs, and control of the argument is returned to Side-2. Side-2 fails to either defeat Side-1's new argument or to bolster its support for the negation of the top claim. Again, Side-1 wins the argument, satisfying the burden of proof of preponderance-of-the-evidence.

Now if we move up to the burden of proof known as dialectical validity, Side-1 must provide one defended path (which it has from "working papers" and "english speaking") and must also defeat any competing/rebutting arguments offered by the other side. Thus, after Side-2 offers its argument with respect to british passport, Side-1's only move is as follows:

```
Find strong support for (AND (BRITISH SUBJECT))
to show that weak evidence (AND (NOT BRITISH PASSPORT))
for claim (AND (NOT BRITISH SUBJECT)) is irrelevant.
Refuted tau would be #:\t230
```

As noted above, Side-1 fails in this attempt and, therefore, loses the argument. The claim can not be established by Side-1 when the proof level is dialectical validity.

The burden of proof of beyond a reasonable doubt is the most difficult proof level to achieve. The input claim must be established with strong, consistent, fully-defended support. As we saw in the dialectical validity example, this is not possible for Side-1 with the claim "British subject" using the given input data and warrants; Side-1 must concede the argument at the first move.

CONCLUSION

In this paper, we have presented a computational model of dialectical argumentation and have demonstrated the roles that burden of proof plays in controlling such argumentation and determining outcomes. We have realized an operational definition of burden of proof as a process that filters possible argument moves and terminates arguments with a declared winner.

Our research simplifies the argumentation context in a number of ways. Perhaps most importantly, there is a mutually agreed upon set of warrants from which to generate argument structures; thus, there is no arguing over the applicability of rules. This is still appropriate for arguing with oneself in the presence of incomplete and competing knowledge. Furthermore, there is no direct appeal to individual cases or the formation of arguments from analogy in our model. These are areas of research that can be incorporated into our current work in the near future, as well as inclusion of a more complete, powerful representation of argument strength.

REFERENCES

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

Loui, R. (1987). Defeat among arguments: a system of defeasible inference. *Computational Intelligence, 3*, 100-106.

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

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

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.

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.

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