Qualitative Argumentation

Arthur Farley

CIS-TR-97-01
January 1997

Department of Computer and Information Science
University of Oregon

Qualitative Argumentation

by

Arthur Farley
Computer and Information Science
University of Oregon
Eugene, Oregon 97403 USA
art@cs.uoregon.edu

Abstract

A basic question for research in model-based, qualitative reasoning is how can we predict or analyze the behavior of complex systems without resorting to completely quantitative models. One difficulty that arises is the ambiguity of results due to conflicting, qualitative indications. Argumentation has long been recognized as a means for resolving issues of belief in situations characterized by incomplete, uncertain, inconsistent, and imprecise knowledge. We explore the application of a model of dialectical argumentation to the domain of qualitative reasoning. Models take the form of qualitative networks, with variables connected by strict or default, positive and negative arcs. A notion of defeat between qualitative arguments, which are represented as paths in a network, is defined. Burden of proof is specified as a flexible means of allocating risk, determining relevant argument moves and eventual outcomes. Examples are presented that illustrate semantics of the approach.

Key Words: qualitative reasoning, dialectical argumentation, burden of proof

Introduction

A basic question addressed by research in model-based qualitative reasoning is how to predict or analyze the behavior of complex systems without resorting to completely quantitative models. One difficulty that arises from the associated loss of precision in qualitative reasoning is the ambiguity of results due to conflicting indications. In a more general setting, argumentation has long been recognized as an appropriate means for resolving issues of belief in situations characterized by incomplete, uncertain, inconsistent, and imprecise knowledge (Polya, 1968). In this paper, we explore the application of a model of dialectical argumentation to issues of ambiguity resolution for the domain of qualitative reasoning about complex systems.

Qualitative Models

Models of qualitative systems and associated argument structures will be specified in a form derived from that used for inheritance reasoning (Horty, 1994). Under this approach, qualitative models of complex systems are represented as *qualitative networks*, being directed graphs of nodes interconnected by typed, directed links. Nodes of a qualitative network, denoted by (names in) small letters, represent the variables and parameters of the system; we will use a small letter from the end of the alphabet (e.g., x, y, z) to represent an arbitrary node of a network.

The nodes of a qualitative network are interconnected by directed links, each link connecting a pair of nodes. There are four link types, corresponding to possible combinations of strength {strict, default} and sign {positive, negative}. The link types =+=> and =-=> denote strict positive and strict negative links, respectively. The meaning of a strict link x=+=>y (or x=-=>y) is that there is a reliable (i.e., always) influence of a change in x upon variable y. A positive link means the change in y is in the same direction as the change in x, while a negative sign indicates that the direction is the opposite. A strict link is used to represent definitional relationships in system models. Link types -+-> and ---> denote default positive and default negative links, respectively. The meaning of a default link x-+->y (or x--->y) is that there is an expected, but perhaps somewhat unreliable, (i.e., usually) influence of a change in x upon the value of variable y. The meanings of the signs are the same as for the strict links. Default links are used to represent observed, but unexplained, regularities or tendencies in system behavior.

Default links have been called "defeasible" (Pollock, 1987), as they can be preempted by stronger, more conclusive indications when reasoning. We use the term "default" for these links to capture the rather strong connection that is intended. We reserve the term "defeasible" for a more general, somewhat weaker strength of influence that exists between nodes of a qualitative network interconnected by paths containing several default links, i.e., the default relation is not transitive. We will discuss the construction of allowable qualitative arguments in the next section.

A qualitative network can serve as basis for answering questions regarding directions of change in model variables given external, input perturbations of parameters or for suggesting changes to parameters that could give rise to a desired (or observed) direction of change in a model variable. Parameters are distinguished from variables in our model of a system in that they have no incoming links within the given qualitative network. Impacts on parameters will be indicated by connecting a special node labeled INPUT to parameters of the network by a set of strict, positive or negative, links. A strict link from INPUT to a parameter indicates a perturbation in the value of the affected parameter in the indicated direction (i.e., positive or negative) as an input. We provide an example of a qualitative network with an associated INPUT in Figure 1. In all figures, strict links will be shown as wider lines.

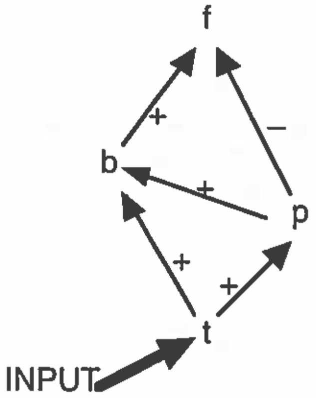

Figure 1

Qualitative Arguments

Given a qualitative network modeling a complex system, we are interested in defining the notion of arguments for and against claims regarding changes in values of variables relative to a given input perturbation. An *qualitative argument* is a directed path in a qualitative network. An argument P(x, π, y) from a *start* node x to a *finish* node y through an intermediate, possibly empty, sequence of nodes π; if node u immediately precedes node v in an argument then u is directly connected to v by a link of the qualitative network. Which directed paths form arguments and with what strengths and signs will be defined by argument construction rules given below.

We characterize an argument in terms of its strength {strict, default, defeasible} and its sign {positive, negative}. We introduce *defeasible* strength for arguments between nodes connected by compound paths involving more than one default link; defeasible arguments capture a qualitative relationship that is weaker than that established by either strict (always) or default (usually) links. If P(x, π, y) is a defeasible positive argument, it represents the idea that "it is reasonable to believe that a change in x will be result in a change in y in the indicated direction". If this path begins at a node INPUT, we are saying that "the given input perturbation could reasonably be considered to result in the corresponding change in y". Distinguishing between strict, default, and defeasible argument strengths will allow us to capture, in a straightforward manner, certain qualitative results that seem intuitively correct. This will be shown through examples presented later in the paper.

Allowable arguments in an inheritance network, with associated strengths and signs, are defined recursively, in the "backward direction" from a given goal node. Links in the network form direct arguments, as defined by the following argument construction rule:

Rule R1: (direct arguments)

- A. given x =+=>y, P(x , ø, y) is a strict positive argument
- B. given x =-=>y, P(x , ø, y) is a strict negative argument
- C. given x -+->y, P(x , ø, y) is a default positive argument
- D. given x --->y, P(x , ø, y) is a default negative argument

In the above rule, x refers to an arbitrary variable or parameter node x or the node INPUT; the symbol "ø" represents the empty sequence of nodes. We extend an argument by adding a new element as start node to its path, creating a *compound argument*. We consider only acyclic compound arguments, defined by the following rule:

Rule R2: (compound arguments)

- A. P(x,π,y) is a strict positive argument not containing z, then
  - (i) given z=+=>x, P'(z,x,π,y) is a strict positive argument
  - (ii) given z-+->x, P'(z,x,π,y) is a default positive argument
  - (i) given z=-=>x, P'(z,x,π,y) is a strict negative argument
  - (ii) given z--->x, P'(z,x,π,y) is a default negative argument
- B. P(x,π,y) is a strict negative argument not containing z, then
  - (i) given z=+=>x, P'(z,x,π,y) is a strict negative argument
  - (ii) given z-+->x, P'(z,x,π,y) is a default negative argument
  - (i) given z=+=>x, P'(z,x,π,y) is a strict positive argument
  - (ii) given z-+->x, P'(z,x,π,y) is a default positive argument
- C. P(x,π,y) is a default positive argument not containing Z (z), then
  - (i) given z=+=>x, P'(z,x,π,y) is a default positive argument
  - (ii) given z-+->x, P'(z,x,π,y) is a defeasible positive argument
  - (i) given z=-=>x, P'(z,x,π,y) is a default negative argument
  - (ii) given z--->x, P'(z,x,π,y) is a defeasible negative argument
- D. P(x,π,y) is a default negative argument not containing Z (z), then
  - (i) given z=+=>x, P'(z,x,π,y) is a default positive argument
  - (ii) given z-+->x, P'(z,x,π,y) is a defeasible positive argument
  - (i) given z=-=>x, P'(z,x,π,y) is a default negative argument
  - (ii) given z--->x, P'(z,x,π,y) is a defeasible negative argument
- E. P(x,π,y) is a defeasible positive argument not containing z, then
  - (i) given z=+=>x, P'(z,x,π,y) is a defeasible positive argument
  - (ii) given z-+->x, P'(z,x,π,y) is a defeasible positive argument
  - (i) given z=-=>x, P'(z,x,π,y) is a defeasible negative argument
  - (ii) given z--->x, P'(z,x,π,y) is a defeasible negative argument
- F. P(x,π,y) is a defeasible negative argument not containing z, then
  - (i) given z=+=>x, P'(z,x,π,y) is a defeasible negative argument
  - (ii) given z-+->x, P'(z,x,π,y) is a defeasible negative argument
  - (i) given z=-=>x, P'(z,x,π,y) is a defeasible positive argument
  - (ii) given z--->x, P'(z,x,π,y) is a defeasible positive argument

In rule R2, z represents either an arbitrary node of the network or INPUT. Throughout the paper, the arbitrary subarguments π or π_i can be empty, i.e., equal ø, unless otherwise noted. The sign of an argument corresponds to the multiplicative influence of signed links along its path in the qualitative network. Of importance is how the strength of an argument is impacted by adding new links. Put simply, an argument path with one default link is of default strength; a path with more than one default link is of defeasible strength. A strict argument contains only strict links.

Given a qualitative network Q, input perturbations as links from node INPUT to system parameters, and a goal variable y, an argument constructed by the rules above is *grounded* iff it starts with INPUT and ends with y, based on the links of Q and the connections of INPUT to Q. For example, given the network in Figure 1, we find three grounded arguments with respect to goal variable f, as follows:

INPUT ~--> f :: P_1(INPUT, t, p, f);
INPUT ~+~> f :: P_2(INPUT, t, p, b, f), P_3(INPUT, t, b, f).

We see that there are both positive and negative arguments regarding the conclusion that variable f increases as parameter t is perturbed upward. This is an example of the ambiguity that often arises in qualitative analyses of complex systems. From an argumentation perspective, we can restate this situation as arguments standing in a conflicting, or contradictory, relationship to one another. Various definitions of conflict relationships between arguments have been given previously (Pollock, 1987; Loui, 1987; Sartor, 1993; Horty, 1994; Farley, 1996). In our terms here, two arguments with the same start and finish nodes, $P(x, \pi 1, y)$ and $P'(x, \pi 2, y)$, *directly conflict* if they differ in sign. More generally, two arguments *conflict* if one directly conflicts with a subargument of the other, i.e., one argument is of the form $(\pi_1, x, \pi_2, y, \pi_3)$, the other is the form (x, π, y), where (x, π, y) and (x, π_2, y) are of opposite sign. In our example above, we have the following pairs of conflicting, grounded arguments: (P_1, P_2) and (P_1, P_3).

Are there methods for resolving qualitative ambiguities, now viewed as conflicts between qualitative arguments? One helpful notion is that of defeat between arguments. Certain pairs of arguments may stand in a stronger contrary relationship than simply being in conflict. In such cases, one argument will be said to defeat the other. Defeat between conflicting arguments is determined by comparing their respective strengths. To make this possible, we state that a strict argument is stronger than a default argument, which in turn is stronger than a defeasible argument.

One argument A *locally defeats* another argument B iff argument B is of the form $(\pi_1, x, \pi_2, y, \pi_3)$, A is of the form (x, π, y), argument A and the subargument (x, π_2, y) are of opposite sign, and A is of greater strength than the subargument (x, π_2, y) of B. In other words, an argument that conflicts with another argument and is of greater strength than the subargument with which it directly conflicts also defeats the other argument. In our example above, the argument $P_4(p, \emptyset, f)$ defeats argument P_2.

Defeating arguments are of course vulnerable to being defeated themselves. In addition, there is a vulnerability associated with intermediate nodes of defeating arguments. If argument $A(\pi_1, x, \pi_2, v, \pi_3)$, where π_1 is a sequence of one or more nodes, is defeated by an argument $B(x, \pi, y)$, where π is a sequence of one or more nodes, then the set of arguments starting at nodes of π_1 and ending at nodes of π that only pass through nodes of A and B is the set of *vulnerable arguments* for the defeat relation between A and B. The vulnerable arguments of a defeat are those that start prior to the beginning of the defeating argument and end at intermediate nodes of the defeating argument. Finally, we have that an argument A *ultimately defeats* an argument B in a qualitative network Q only if A locally defeats B, A is not defeated, and none of the associated vulnerable arguments are defeated.

Burden of Proof

A *qualitative claim* is a statement regarding the influence of one element (i.e., variable, parameter, or INPUT perturbation) of a qualitative network on another variable of the network. A positive claim will be denoted as x~+~>y, while x~-~>y will denote the negative claim. Each grounded argument P(INPUT, π, y) supports one of these two claims regarding the direction of influence of a specified input perturbation on variable y of the model.

To determine the acceptability of a qualitative claim, we must decide which type of error, either of commission (i.e., false positive acceptance of a claim) or omission (i.e., false negative rejection of a claim), we are more willing to have occur. We would like to adjust this allocation of risk depending upon estimated consequences of wrongly accepting or rejecting a qualitative claim between a particular model element and given variable.

We provide the ability to allocate risk in a flexible manner by specifying a burden of proof. A burden of proof is a parameter to the argument process, not a property of the reasoning system. There are two aspects to the specification of a burden of proof: (i) which side of a claim (positive or negative) bears the burden and (ii) what level of proof is required for acceptance of the claim. The first aspect addresses whether one is more concerned about accepting false positives, where the burden is placed on the claim of interest, or false negatives, where the burden is placed on the claim of opposite sign. Sometimes, we want a good argument supporting a claim before accepting it, where the risk associated with wrongly accepting the claim is perceived to be high. On other occasions, where accepting a claim has potentially high value and little perceived risk, we may demand a good argument against the claim (i.e., of opposite sign) before denying its acceptance.

One domain in which the notion of burden of proof has long been applied is that of legal reasoning. A different burden of proof may be mandated at each stage of a legal process or for a different type of legal action. For example, arguments sufficient to indict someone need not be as convincing as those needed to convict someone of a criminal offense. When considering conviction in criminal cases, we tend to be concerned about errors of commission (i.e., a finding of guilt when not guilty) and, thus, place a high burden of proof on the side arguing for guilt.

The second aspect of burden of proof, that of proof level, addresses the issue of what is considered to be a good, or sufficient, argument. Proof level will be based upon the following notions of defendable and justifiable arguments. A *defendable argument* is an argument that is not ultimately defeated by any other argument of a given qualitative network and INPUT perturbation. A *justifiable argument* is a defendable argument with the added condition that every argument that conflicts with it is ultimately defeated.

We now define the following three proof levels:

- *scintilla of evidence (se)*: there exists a defendable argument supporting the claim;

- *preponderance of evidence (pe)*: there exist more defendable arguments supporting the claim than its negation;
- *dialectical validity (dv)*: there exists a justifiable argument supporting the claim.

Winning a scintilla of evidence argument for a claim merely requires that some argument for the claim be defended against all attacks. In a dialectical context, this means that only defeating arguments can be considered by the side opposing the claim; a conflicting argument of lesser or equal strength is irrelevant. Preponderance of evidence requires a means for assessing relative strengths of sets of conflicting, defendable arguments. We assume that having a greater number of defendable arguments for a claim represents a stronger case. Under preponderance of evidence, if the opposing side finds an equally strong argument, this can not be ignored; the argument must be counteracted, either by defeating it or by finding a new argument of equal strength. Finally, dialectical validity requires not only that some argument be defendable but also that any conflicting argument be defeated, even though that argument is not any greater in strength.

We can now define the semantics associated with a given qualitative network Q and INPUT perturbation in terms of the sets of claims acceptable under differing burdens of proof. We denote by *D(Q, INPUT)* the set of defendable arguments, given network Q and perturbation INPUT. Similarly, we denote by *J(Q, INPUT)* the set of justifiable arguments. By our above definitions, D(Q, INPUT) is a subset of J(Q, INPUT). We denote by *C(Q, INPUT, L)* the set of claims that are acceptable with proof level L, given Q and INPUT. The set C(Q, INPUT, L) is derived from sets D(Q, INPUT) and J(Q, INPUT), as per our definitions above. A claim *c* is an element of C(Q, INPUT, *se*) iff there exists an argument for *c* as claim in D(Q, INPUT). A claim *c* is an element of C(Q, INPUT, *pe*) iff there exists more arguments for claim *c* in D(Q, INPUT) than there are arguments for the complement of *c* in D(Q, INPUT). A claim *c* is an element of C(Q, INPUT, *dv*) iff there exists an argument for claim c in J(Q, INPUT).

The three proof levels defined above create a hierarchy of acceptable claims based on set inclusion. For any Q and INPUT, C(Q, INPUT, se) contains C(Q, INPUT, pe), which contains C(Q, INPUT, dv). Only C(Q, INPUT, se) can contain contradictory claims, i.e., claims between the same variables with opposite signs.

Examples

We now turn our attention to a number of examples that demonstrate general principles of our approach and illustrate the impact that burden of proof has upon argumentation semantics. In the example from Figure 1, both positive and negative claims can win only scintilla of evidence arguments; there is exactly one defendable, defeasible argument supporting each claim.

In Figure 2, we see one positive argument P_1(INPUT, e, d, b, a) of defeasible strength for a positive input perturbation to e resulting in a positive change in a. However, this argument is defeated by argument P_2(d, c, a), a negative argument of default strength defeating subargument P_3(d, b, a) of P_1. As a result, the only grounded, defendable argument is P_4(INPUT, e, d, c, a). While it is only of defeasible strength, it is not attacked by any other defendable argument. Thus, the negative claim INPUT~--->a can win arguments for all three burdens of proof.

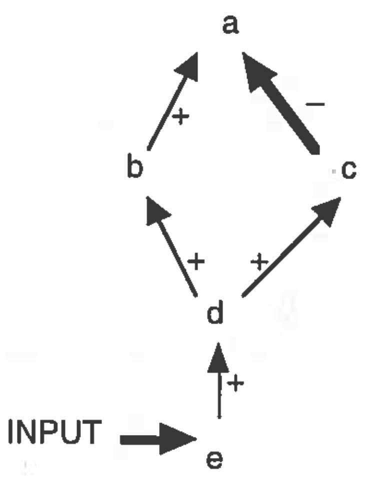

Figure 2

We make the situation slightly more complex by adding parameter f connected to e and letting e have a direct, negative, default impact on c, as shown below in Figure 3. As before, the defeasible, positive argument P_1(INPUT, f, e, d, b, a) is defeated by the default, negative argument P_2(d, c, a). This time, however, the defeasible negative argument P_3(INPUT, f, e, d, c, a) is defeated by the default negative argument P_4(e, ø, c), which conflicts with the positive default subargument P_5(e, d, c) of P_3. By defeating P_5, argument P_4 reinstates P_1; it defeats a vulnerable argument associated with the defeat of P_1 by P_2. The only other defendable, grounded argument in this example is P_6(f, 3, c, a), which is a defeasible, positive argument.

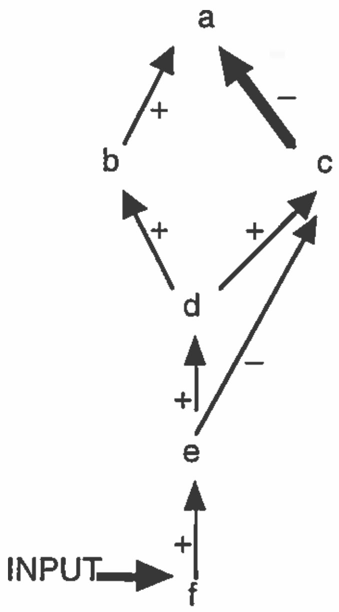

Figure 3.

We see that for the network of Figure 3, the claim INPUT~+~>a can win arguments (i.e., is acceptable) up through a burden of proof of dialectical validity. In the previous two examples, if we did not consider the structure of the arguments involved and their conflicting interactions, the qualitative indications would be ambiguous. By considering defeat relations among arguments, we see there is a clear prediction for direction of influence according to our argumentation semantics.

We can create more ambiguous situations, such as the one presented in Figure 4. Here no arguments are defeated. As a result, there are three grounded, defendable arguments: P_1(INPUT, e, b, a), P_2(INPUT, e, c, a) and P_3(INPUT, e, d, a). Argument P_1 indicates a negative influence, while P_2 and P_3 indicate positive. The claim INPUT~-~>a can win arguments with a burden of proof of scintilla of evidence. However, by being able to outweigh the set of arguments available for the negative claim, the claim INPUT~+~>a can win arguments up through a burden of proof of preponderance of evidence. Recall that scintilla of evidence is the only burden of proof for which both positive and negative claims can be considered acceptable. Use of scintilla of evidence is appropriate for situations where there is little perceived risk in making errors of commission. Scintilla of evidence allows acceptance of a claim even when the other side has more arguments in its favor, as long as there exists a defendable argument for the claim.

If any arc in Figure 4 were made strict, the argument containing it would become dominant. The associated claim would win arguments up through dialectical validity.

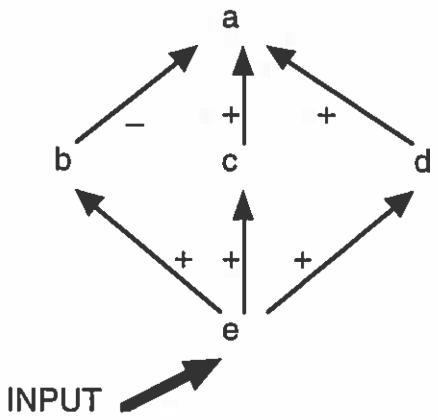

Figure 4

Argument Process

Now that we have defined a structure for arguments as acyclic paths in a qualitative network and characterized several important properties of and relationships between arguments, we describe a process whereby we can decide whether to accept a claim or not. The decision will be made through a process of dialectical argumentation under a given burden of proof. The process model presented here is based upon a general model of dialectical argumentation described earlier (Freeman, 1993; Farley and Freeman, 1995), here modified and adapted to the circumstances of qualitative argumentation.

A dialectical argument has two sides, where *Side-1* argues in favor of an input claim and *Side-2* against the claim, i.e., in support of its negation. The argument process begins with Side-1 attempting to find a grounded argument for the input claim. Search for a supporting argument proceeds from the goal node toward the INPUT node in a backward-directed manner, according to the argument construction rules given above. If no support can be found, the argument ends with a loss for Side-1; the input claim is not accepted, as all burdens of proof require Side-1 to find at least one grounded, defendable argument. Subsequently, the two sides alternate as *active side* of the argument. A side remains active until it succeeds in creating a check condition or runs out of moves and must concede the argument. A *check condition* for a side S is a situation such that, if the other side can not successfully respond, side S wins the argument. Except for the initial situation, when Side-1 must generate a grounded argument for the claim, the active side is faced with a set of *check arguments* for the other side, i.e., arguments responsible for the other side holding the check condition.

When active, a side selects an argument move to apply from a set of possible moves. For this, a side can apply one of two primitive functions that search the qualitative network for relevant arguments. The first is *find-arguments (x, g, s, Q)* which searches for argument paths from node x to node g of sign s in qualitative network Q. The function returns a list of argument paths in order of decreasing strength, or it returns an empty list indicating no such paths exist. The other function, *find-conflicting-arguments(A, Q)*, finds arguments that conflict with argument A in network Q. It is equivalent to the union of directly conflicting arguments for each pair of nodes in A. This function is implemented by calls to the function *find-arguments* with parameters being pairs of node from argument A and the complement of the sign of the subargument in argument A. Defeating and rebutting (i.e., not defeating) conflicts can be recognized in the process.

Whether an argument is adequate to generate a check condition for the active side depends upon the burden of proof specified. For example, under a burden of proof of dialectical validity, Side-2 can consider both defeating and rebutting arguments in response to Side-1's arguments. If Side-2 finds a conflicting argument, Side-1 must defeat Side-2's response or propose a completely new argument for its claim; otherwise, it must concede the argument. Side-2 can continue posing counterarguments to a given argument, all of which Side-1 must defeat if it is to prevail under a burden of dialectical validity.

If the burden of proof is preponderance of evidence, then Side-2 must generate a counterargument of strength equal to that proposed by Side-1. If it can do this, Side-1 finding another argument in favor of the claim is sufficient to regain a check condition. As long as Side-1 can come up with more arguments of strength greater than or equal to those that Side-2 produces, it will prevail. If the burden of proof is merely scintilla of evidence, Side-2 can only consider defeating arguments in response to a Side-1 argument. Side-1 need not defeat rebuttals to win the argument under this burden of proof; it must merely defend some argument against defeat; if Side-2 can defeat an argument, Side-1 can abandon it in favor of another supporting the input claim.

We can characterize burden of proof in terms of where it places the "burden of defeat", i.e., which side must defeat the other's arguments. In the case of scintilla of evidence, the burden of defeat is on Side-2; while under a burden of proof of dialectical validity, the burden of defeat is on Side-1. Under preponderance of evidence, neither side takes on the burden of defeat. This is a free for all, where piling up more arguments of greater or equal strength in one's favor is sufficient. Overall, we see that burden of proof plays several roles in the process of dialectical argumentation: (i) as basis for deciding the relevance of arguments found by the active side; (ii) as basis for deciding the sufficiency of the outcome of an active side's move (i.e., whether a check condition has been realized); (iii) as basis for determining that an argument is over; and (iv) as basis for determining whether a claim is accepted or not.

Conclusion

This paper reports results of an initial study applying notions developed previously in the field of argumentation to issues arising in qualitative reasoning about complex systems. The domain of qualitative reasoning is well-suited for the application of argumentation, as qualitative reasoning contexts can be characterized as dependent upon incomplete, imprecise, uncertain, conflicting indications. Dialectical argumentation provides a method for considering both sides of a claim and making a decision as to its acceptability based upon an allocation of risk .

We have previously applied our model of argumentation to domains of legal reasoning and nonmonotonic inheritance (Farley and Freeman, 1995; Freeman and Farley, 1996; Farley, 1996). The current effort is most closely related to the latter work on inheritance reasoning, where we considered inheritance arguments as paths in inheritance networks. The semantics of links in an inheritance network and the set of allowable arguments differ somewhat from those defined here for qualitative networks, but the definitions of defeat and burden of proof are adopted directly.

We have modified our Scheme implementation that was previously completed for the inheritance reasoning domain to perform the qualitative argumentation outlined above. This implementation determines the sets of claims acceptable for each of the three burdens of proof by first computing the sets of justifiable and defendable arguments. Next steps in our research will be to reimplement the system in terms of the dialectical argument process described above and to explore problem solving and diagnosis in terms of more complex qualitative networks that model interesting, real-world systems.

References

Dung, P.M. 1995. “On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games”, *Artificial Intelligence*, 77(2), 321-358.

Farley, A.M. 1996. “Dialectical Nonmonotonic Inheritance”, UO-CIS-TR-96-15, submitted to *Computational Intelligence*.

Farley, A.M. and Freeman, K. 1995. “Burden of proof in legal argumentation”, in *Proceedings of Fifth International Conference on Artificial Intelligence and Law*, 156-163.

Freeman, K. 1993. *A Computational Model of Dialectical Argumentation*, Ph.D. Thesis, Computer and Information Science Department, University of Oregon.

Freeman, K. and Farley, A.M. 1996. "A computational model of argumentation and its application to legal reasoning", to appear in *Artificial Intelligence and Law.*

Horty, J.F. 1994. "Some direct theories of nonmonotonic inheritance", in Gabbay, D.M., Hogger, C.J, and Robinson, J.A. (eds.), *Handbook of Logic in Artificial Intelligence and Logic Programming*, Oxford Press : New York, 111-188.

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

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

Polya, G. 1968. *Mathematics and Plausible Reasoning* (2nd ed.), Princeton University Press: Princeton, NJ.

Prakken, H. 1993. "An argumentation framework in default logic", Annals of Mathematics and Artificial Intelligence (9), 93-132.

Sartor, G. 1993. "A simple model for nonmonotonic and adversarial legal reasoning", in *Proceedings of Fourth International Conference on Artificial Intelligence and Law*, 192-201.