Argumentative Reasoning Theory: Explanation Aware Knowledge Representation

Andrew Potter

Sentar, Inc.
4900 University Square, Suite 8,
Huntsville, Alabama 35816 USA
apotter@Sentar.com

Abstract

Argumentative Reasoning Theory (ART) is a theory of knowledge representation, reasoning, explanation, and argument interaction. It is designed to support intelligent human-computer collaboration. ART provides the ability to represent reasoning in a form that is computable, intuitive, and amenable to discovery. By integrating Toulmin's model of argumentation, Mann and Thompson's Rhetorical Structure Theory, and Perelman and Olbrechts-Tyteca's strategic forms of associative and dissociative reasoning, ART defines an ontology for representing and manipulating of argument structures. Arguments, when satisfied, are instantiated into a dynamic rhetorical network that represents the system's model of the situation. Two modalities of instantiation are used: inferential instantiation is used when the claim is inferred from the ground; synthetic instantiation is used for descriptive argumentation where both ground and claim must be satisfied for the argument to be instantiated. The instantiation process maps arguments into the network using interaction links. Links are defined for a range interactions, including accrual, concomitance, backing, substantiation, dissociation, rebuttal, undercut, and confusion. Interaction detection may be accomplished using logical, ontological, and inventive-dissociative detection. Knowledge discovery is supported through logical and analogical means. Through ontologically normalized representation of argumentative knowledge, it becomes possible to detect the opportunity for analogical discovery.

Introduction

The notion that argumentation theory could be used to motivate the development of technologies for intelligent human-computer collaboration has been explored by numerous researchers. Among these, Ye (1995) and Ye and Johnson (1995) investigated expert system interaction with human users. They found that a system capable of presenting arguments persuasively is more likely to be regarded as a credible resource for resolving complex issues. Moulin, et al. (2002) argued that argumentative reasoning strategies could be used to make agents more persuasive and proposed that Perelman and Olbrechts-Tyteca's (1969) analysis of argumentation could be used in this endeavor. Along similar lines Grasso (2002) used rhetorical schemas for modeling argumentative dialogues, with the objective of providing participants with a familiar behavioral model. Wærn and Ramberg (2004) proposed a system that would use Mann and Thompson's (1988) Rhetorical Structure Theory (RST) to construct explanation networks and Toulmin's (1958) model to find paths through these networks, resulting in a dual level knowledge system that would support both inferencing and explanation derivation.

Clearly, if the human propensity for argumentation could be imparted to computers, computers could in turn be used to engage humans in argumentation and in the complex problem-solving processes enacted through argumentation. And yet what has not emerged from previous studies is a general theory of reasoning for use in human-computer collaboration. For humans and computers to collaborate, they must reason together, and in order to reason together, they must share common ground in rhetoric and argumentation. Establishing that common ground is the objective of the approach described here, called Argumentative Reasoning Theory (ART).

ART is a theory of rhetorical reasoning and knowledge representation. It is intended to support intelligent human-computer collaboration under conditions involving uncertainty and change. It provides the ability to represent complex, dynamic argumentative structures in a way that is machine-processable, humanly intuitive, and amenable to discovery of new argumentative structures. The theory is based on Toulmin's model of argumentation (Toulmin 1958), Mann and Thompson's Rhetorical Structure Theory (RST) (Mann & Thompson 1988), and Perelman and Olbrechts-Tyteca's strategic forms of associative and dissociative reasoning (Perelman & Olbrechts-Tyteca 1969). The Toulmin model provides the framework for an ontological specification of argumentative information. RST adds to this the ability to express interrelationships among argument components in meaningful and useful ways. RST also provides a conceptual apparatus for construction of dynamic rhetorical networks of interrelated arguments. Associative and dissociative reasoning leads to the development of conceptual tools for defining and carrying out argumentative strategies. Because ART is

April, 2006
Revised April, 2007

modeled on natural discourse, rhetoric, and argumentation, the resulting knowledge structures and processes possess salient characteristics of intelligent interaction.

Using event-driven or goal-seeking scenarios, arguments may be continuously evaluated in response to the dynamics of the external situation. When satisfied the argument is instantiated and mapped into a coherent rhetorical network that models the situation. This network provides the basis for reasoning operations, and the resulting network is itself a rhetorical explanation of the system's reasoning. It may be filtered, summarized, and rendered as appropriate for human-computer interaction. Because the network conforms to a rhetorical model of representation, the technologies used for uncertainty management, defeasibility, conflict resolution, and discovery may be defined as analogues to comparable human judgment.

Theoretical Foundations

Three strands of research provide the basis for the theory. These are the Toulmin model (Toulmin 1958), Rhetorical Structure Theory (RST) (Mann & Thompson 1988), and Perelman and Olbrechts-Tyteca's strategic forms of argumentative processes (Perelman & Olbrechts-Tyteca 1969). The Toulmin model is a theory of argumentation. As shown in Figure 1, Toulmin defined an argument as consisting of six elements: a *claim*, a *ground*, a *warrant*, a *backing*, a *qualifier*, and a *rebuttal*. The *claim* is what the argument purports to demonstrate. The *ground* is the data that support the claim. The *warrant* establishes the linkage between ground and claim. The *backing* is a policy, law, argument, or fact that substantiates the warrant. The *qualifier* is an indication of the strength of the argument. The rebuttal is any counter-argument that might refute the argument.

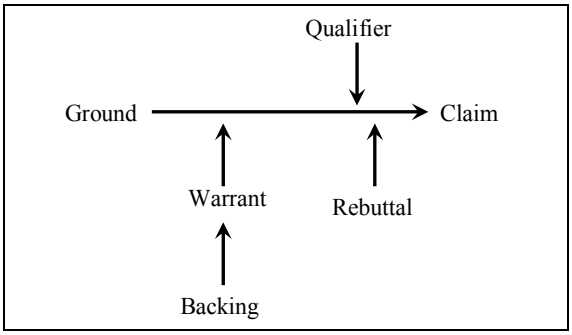

Figure 1: Toulmin Model

Rhetorical Structure Theory (RST) is a theory of text coherence. RST defines the coherence of a text in terms of the way its parts, or *text-spans*, relate to one another. It postulates a small number of schemas for defining the possible structural relationships among spans and defines a set of rhetorical relations that may be used when applying a schema to a set of text spans. An RST analysis of a coherent document defines a hierarchical structure representing the rhetorical interrelationships of the spans comprising the document. Figure 2 shows an example of an RST structure using the EVIDENCE relation.

Perelman and Olbrechts-Tyteca's strategic forms of argumentative processes include *associative* and *dissociative* forms. *Associative* forms are usually connective. Connections are established when, to use Toulmin's terms, a claim is supported by ground and warrant. However, associative forms may also be disconnective. A disconnective form is used in a rebuttal which denies a claim by challenging the grounds. It breaks an argumentative link. *Dissociative* forms are not merely disconnective; they are disruptive, as they may challenge the underlying theory or backing that gives credence to the argument. Dissociative arguments do not merely rearrange the links; they alter the nature of the argument.

Figure 2: RST Diagram of the EVIDENCE Relation

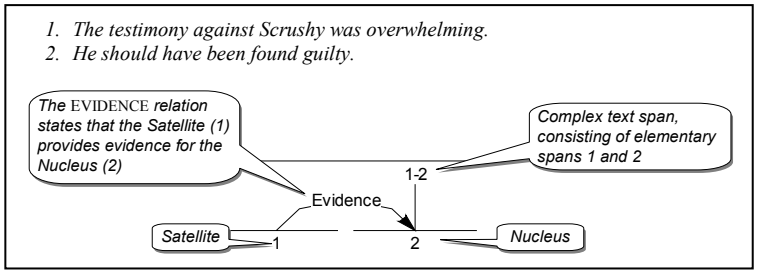

Argumentative Reasoning Theory

The Toulmin model and Rhetorical Structure Theory share a common thread with classic rule-based representations. They are each concerned with plausible relationships between portions of information. Rules define dependencies between conditions and consequences. Toulmin arguments use warrants to link grounds with claims. RST identifies functional relationships between satellites and nuclei.

Rules define a relationship between two situations, one being the condition and the other the consequent. The situation described in the consequent is dependent upon the situation described in the condition. The condition specifies what is necessary and sufficient for the consequent to be asserted or performed. In stating that one situation is contingent on another, the rule says nothing about why this is so. As such it offers little in the way of explanation. The rule is convincing only to the extent that its consumer has some familiarity with the domain environment hosting the rule.

Toulmin warrants are comparable to rules (Sartor 1993). The ground is the condition, and the claim is the consequent. But Toulmin makes provision for information not available using rules. The qualifier distinguishes the relationship between ground and claim with respect to certainty. The backing may be used to offer an assurance of the acceptability of the warrant. And the rebuttal provides the option of specifying possible counter-arguments or claims.

From the perspective of Rhetorical Structure Theory, the Toulmin warrant corresponds to a rhetorical relation. This can be seen in RST relations where the satellite gives evidence or explanation for a nucleus, as with the EVIDENCE and VOLITIONAL-CAUSE relations. Applications of these relations indicate that the nucleus relies on the satellite for support, just as with rules the consequent requires the support of the condition, and with Toulmin arguments the grounds are sufficient to assert, as delimited by the qualifier, the claim. As will be developed in this paper, however, the mapping of RST relations into argumentative formulations is not limited to a select group of relations.

Toulmin and RST provide more expressive models of the situation than can be specified using rules. The Toulmin model captures everything that the rule does but also offers the ability to include backing, qualifier, and rebuttal. RST provides the means to specify the rhetorical relationship between one text-span and another, and it provides a means for expressing a much richer array of inter-propositional relationships than possible through either rules or Toulmin arguments. RST provides a model for representing a complex structure consisting of multiple interlocking contingent relationships. This presents the possibility for dynamic construction of highly expressive knowledge networks.

An Argument Ontology

The ontology presented here defines a conceptualization for representing argumentative knowledge above the sentence level. As a theory of how the constituents of an argument interrelate, it unifies the Toulmin model with Rhetorical Structure Theory. Generally, Toulmin's grounds and claims correspond with RST satellites and nuclei. To this extent, RST supplies relations for describing the nature of the relationship between ground and claim. However, argumentative discourse goes beyond inferential relations to include a variety of structures useful in argumentative discourse. Directly argumentative structures such as CONCESSION, ANTITHESIS, and JUSTIFY play important roles (Azar 1999), and ultimately the full range of rhetorical relations may be needed if the knowledge representation model is to support human-computer collaboration. To realize this, it is necessary to formalize these concepts ontologically. By representing discourse ontologically, arguments thus specified are not merely a means of affirming claims on the basis of grounds; they are objects of knowledge and may be treated accordingly.

As shown in Figure 3, an *Argument* defines a *Warrant* and a set of *Interactions*. The Warrant defines the *Nucleus* (Claim) and the *Satellite* (Ground). The Nucleus is an ontologically normalized statement, usually of a domain specific nature. The Satellite specifies the satellite statement and its *Relation* to the Nucleus. The Relation identifies the RST relation and characterizes its modality as either synthetic or inferential (modality will be discussed in the next section).

The Argument also identifies a *qualifier, qualification ratio,* and argument *Interactions.* The qualifier is a static value indicating the level of certainty of an individual argument, and may be either *conclusive* or *supportive*. The qualification ration is dynamic and is defined by the interactions in play between the an instantiation of the argument and other argument instantiations.

Interactions define the possible relations an instantiated argument may have with some other argument. In the Rhetorical Network, various Arguments may be linked together by means of such Interactions. For example, when the Nucleus of one Argument unifies with the Satellite of another, the *substantiation* interaction is specified; when the Warrant of an argument appears as the Nucleus of another, the *backing* interaction is used; and when two Arguments share the same Claim, the *accrual* interaction is used.

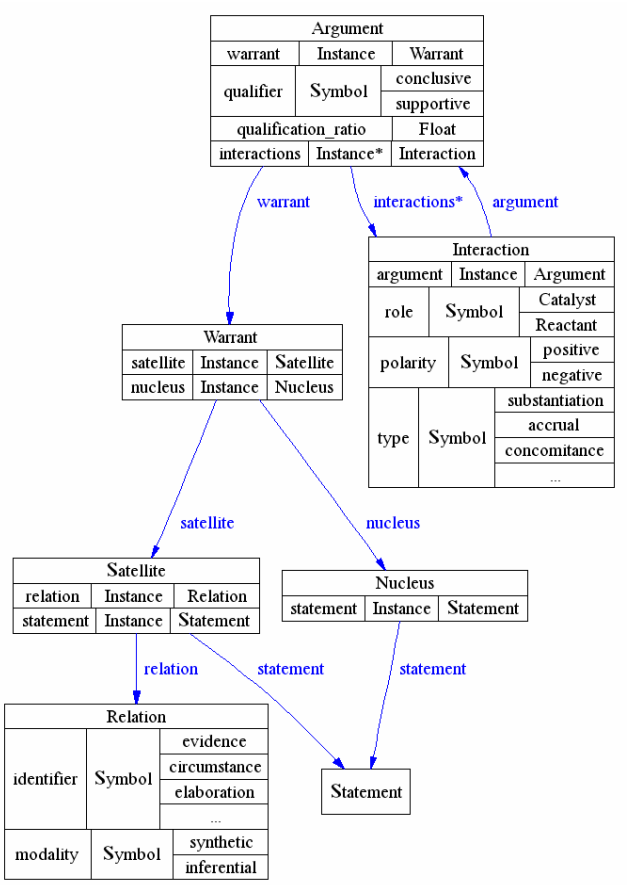

Figure 3: An Argument Ontology

Rhetorical Networks

Arguments are generally thought of as consisting of premises and conclusions, such that for the conclusion to be accepted, adequate premises must be provided to support it (Juthe 2005). Rule-based representations are consistent with this, to the extent that the conditions play the part of premises and the consequent is the conclusion. Following this view, a rule is triggered when its condition

is satisfied, resulting in the consequent being asserted. While definitions of this sort may be sufficient for use in logical and quasi-logical inferencing, argumentative reasoning requires a more encompassing approach. Argumentative discourse seldom follows a simple premise-conclusion linkage. This can be seen in the following antithetical argument:

> *Rather than waste time attending classes, Alan bought his diploma on the Internet.*

The example is argumentative to the extent that the ground (wasting time attending classes) is intended to increase the reader's positive regard for the situation presented in the claim (buying a diploma on the Internet). In asserting the relevance of the ground to the claim, there need be no implied generalization that anyone who avoids attending classes is likely to buy a diploma on the Internet. For this argument to be instantiated, both the ground and the claim would need to be satisfied. The force of the argument is not that one part is used to establish the other, but rather that the parts occur in a significant relationship to one another—that relationship in this case being antithetical.

Further, to understand a situation is not simply to comprehend a collection of discrete facts and inferences, but rather to realize how they combine to produce an integrated view within the scope of a universe of discourse. Therefore the approach calls not for asserting claims on the basis of grounds, but for instantiating arguments as constituting specific applications of the argumentative form. Moreover, there are two modalities of instantiation. One modality applies to arguments wherein the claim may be inferred from the ground; this called *inferential instantiation*. The second is applicable to arguments where both ground and claim must be satisfied for the argument to be instantiated. This is called *synthetic instantiation*. EVIDENCE is inferential; ELABORATION is synthetic.

When an argument is satisfied, either inferentially or synthetically, it is instantiated into a rhetorical network. The rhetorical network is a unified expression of the world as known to the system. These networks are similar to the inference networks used by Pollock (1995) and the explanation paths proposed by Wærn and Ramberg (2004). When an argument is instantiated, any variables specified in the ground and claim are bound, and the argument is mapped into the network of instantiated arguments. The network links are specified in terms of the argumentative interactions shared among the arguments.

Argument Interaction

A rhetorical network is a map of the interactions among arguments. In a monotonic argumentative discourse, interactions are generally positive, consisting of substantiation and backing. That is, arguments provide support for other arguments, either by affirming their grounds or by supplying backing for their warrants; the resulting linkage leads eventually to the final claims of the discourse. However, in an argumentative discussion involving multiple agents, the patterns of interaction are unlikely to be as simple as this. Arguments may conflict with one another by disputing each other's claims, their grounds, or their warrants. They may converge on a single claim, or, from a single ground, multiple claims may issue forth. Thus, the first consideration is that interactions must be thought of as having a *type*. There are eight types of argument interaction, as summarized in Table 1.

Table 1: Summary of Argumentative Interactions

Substantiation	The claim of one instantiated argument is the ground of another
Rebuttal	The claim of one argument disputes that of another
Backing	The claim of one instantiated argument substantiates the warrant of another
Undercut	The claim of one instantiated argument disputes the ground of another
Dissociation	The claim of one instantiated argument disputes the warrant of another
Accrual	Two instantiated arguments share the same claim, but their grounds differ
Concomitance	Two instantiated arguments share the same ground, but their claims differ
Confusion	The instantiated ground of one argument is incompatible with that of another

Before discussing these interaction types in detail, it is useful to introduce some other key concepts involved in argument interaction. These include *locus*, *polarity*, *catalyst*, and *reactant*. The *locus* of interaction identifies the elements of the arguments denoted in the interaction. Some interactions denote claims, some grounds, some combinations grounds and claims, and some combinations of claims and warrants. To provide a clear understanding of the nature of an interaction type, it is necessary to locate these elements precisely. Hence every type has a locus.

Interactions are typically, but not always, *catalytic*. That is, one argument attempts to influence another, but the favor is not returned. For example, one argument that substantiates another is not itself substantiated by the latter unless the reasoning is circular. Thus, in an interaction, one argument may be designated as the *catalyst* and the other as the *reactant*. If the influence exerted by the catalyst is supportive, the *polarity* is *positive*. If the influence is resistant, the *polarity* is *negative*. In some interactions, the influence may be benign, or *neutral*. When one argument substantiates another, the catalyst exerts a *positive* polarity on the reactant. When an argument rebuts another, the catalyst exerts a *negative* polarity on the reactant. With these concepts in mind, it now becomes possible to explore the details of the various interaction types.

Substantiation and Rebuttal

Substantiation occurs when the claim of one instantiated argument is used as the ground of another. In the following example, the first argument substantiates the grounds of the second:

1. *According to library records, Jack failed to return the book on time*
2. *Jack failed to return the book on time, so he should pay a fine*

The claim of the first argument matches the ground of the second. Thus the locus of interaction is the unit *Jack failed to return the book on time*. Since the first argument provides support for the claim made in the second, the polarity is positive, and the catalyst is the first, and the reactant is in the second.

In rebuttal, the locus of interaction is on the claims made by the interacting arguments. One claim disputes the other, for example:

3. *Jack failed to return the book on time, so he should pay a fine*
4. *Jack already paid, so he owes nothing*

The polarity of Rebuttal is negative, the catalyst is in the rebutting argument, and the reactant is the argument subjected to rebuttal. Thus rebuttal is a process of claim and counterclaim.

Backing, Undercut, and Dissociation

In Toulmin theory, Backing is the policy, law, argument, or fact that substantiates the warrant. More generally, *Backing* is any argument that substantiates a warrant. Thus it is an argument with positive polarity whose loci of interaction resides in the claim of the catalyst and the warrant of the reactant.

Some researchers use the term *Undercut* to refer to a claim that challenges a warrant (Pollock, 1995; Prakken, 2005). A claim that challenges a warrant challenges the standing of the argument itself; that is, if successful, there is no argument—the ground is no longer a ground and the claim is not longer a claim; rather they are henceforth dissociated units. The term *Undercut*, as used here refers to a less disruptive form of challenge, one in which the catalyst challenges the ground of the reactant; thus it undercuts the claim. The more severe form of challenge, where a claim challenges a warrant, is referred to as *Dissociation*. Whereas an Undercut challenges a particular instantiation of an argument, Dissociation extends to the full range of instantiations:

5. *Charging overdue fines is an unfair practice*

This challenge is not specific to Jack's dilatory returns, but extends to anyone who has failed to return a book on time.

Accrual, Concomitance, and Confusion

The basic idea of accrual is that the whole is greater than the sum of the parts. Multiple arguments leading to the same claim, or multiple instantiations of the same argument leading to the same claim, might collectively strengthen the claim. There is a strong cultural disposition toward this expectation: getting a second opinion before undergoing a risky medical procedure seems prudent; corroborating testimony to establish evidence in a legal proceeding may influence juries; if someone in addition to yourself saw a flying saucer in your backyard last night, you may be able to avoid time in a padded cell; the scientific method emphasizes the notion of repeatability. The very idea of democracy is based on gaining a plurality of opinion. To be a rebel just like everyone else is an essential phase in coming of age in post-modern society. Anytime two travelers reach the same destination, whether they each followed the same path, or by differing paths so arrived, accrual is not far away.

There seem to be two basic forms of accrual. These are *repeatability* and *convergence*. With *repeatability* the mere fact that there is a multiplicity of instances contributes to the persuasiveness of the claim. Repeatability occurs when the warrants of the accruing arguments are the same. These involve individual instantiations of the same argument. Repeatability includes a range of argumentative manifestations, including corroborating testimony, second opinions, scientific observations, polling, gossip, and riots. The value and necessity of this form of accrual depends on its particular variety and the details of the domain.

With *convergence* the warrants differ but the claims are the same. Prakken's (2005) jogging example uses convergent accrual, where one argument uses rain as a reason not to go jogging and the other uses heat as an excuse, with the result that the two arguments converge upon a single claim. Using this example Prakken demonstrates that accrual may under some circumstances be weakening—the combination of rain and heat could be pleasant conditions for jogging. It may be that accrual weakening is confined to convergence. Of the varieties of repeatability mentioned—testimony, second opinions, scientific observations, sampling, gossip, and riots—it would seem that repeatability is likely to be strengthening. Obviously the degree of strengthening will vary not only on the basis of variety, but from one domain to another, and from one instance to another. This seems consistent in a general way with Pollock's (1995) view that accrual is not a matter of presuming the whole to be greater than the

sum of its parts—it is necessary to contribute to the mix some additive attainable only by conjoining the elements.

Concomitance occurs when two arguments use the same ground to establish distinct claims. A well-behaved concomitance is non-catalytic, and the polarity is neutral. However there may be incompatibility between the claims, as in the following example:

6. *When I am hungry, I can eat a whole pie.*
7. *When I am hungry, I can eat a whole cake.*

Eating the cake or eat the pie presents no problem, but eating both may prove a challenge. Depending of the rigor of the ontology, such incompatibilities might be detectable, but otherwise concomitance should be regarded with suspicion.

Confusion occurs when incompatible grounds are instantiated. If either of the ground is also the claim of some other argument, the condition may more appropriately be handled as an undercut, where one argument disputes the ground of another. However, if neither of the grounds are thus substantiated, then confusion results. This can occur when the grounds are provided as input from some other agent. The locus of interaction is in the grounds, the polarity is negative, and both arguments are reactants.

Effects of Interaction

Detecting and describing argument interaction is useful in understanding networks of argumentation, but only insofar as this understanding supports evaluation of the quality of the claims. Each argument ground has a qualifier, defined as either conclusive of supportive, that indicates the certitude of its relation to the claim. This in itself provides only a partial view, however. An argument, once instantiated, interacts with other arguments in the rhetorical network in ways that affect its certitude. Depending on the type of interaction, one argument may strengthen or weaken another. A positive polarity will strengthen the argument; a negative polarity will weaken it.

The approach described here is based on Fox and Das (2000). The strength of a claim is measured as a *qualification ratio*. The qualification ratio is a measure of the certitude of an argumentative structure. This ratio is derived from the qualifiers associated with the interacting arguments. It may be found by determining the ratio of conclusive positive interactions to total positive interactions and the ratio of conclusive negative interactions to total negative interactions, and then subtracting the latter from the former:

$$QR = TCP / TP - TCN / TN$$

A higher ratio indicates a higher level of certainty. Given two competing arguments, the one with the higher qualification ratio is more certain. Note however that this does not lead to retraction of the weaker argument. The rhetorical network accommodates the presence of defeated structures, as well as their restoration to grace when the situation changes. Retraction of a weaker argument, or of any element that has become obsolete, would be undertaken on extra-theoretical grounds, dependent upon the needs of a particular application. To this extent, the reasoning model adopts an open world assumption.

Explanation

The reasoning processes performed by a reasoning system must be explainable, otherwise these processes would not, by definition, be reasonable. The criteria used by the system to substantiate its claims and those presented to the user must be identical. If there are deficiencies in the arguments advanced by the system, or in the system's ability to represent the arguments, these defects should be clearly identified rather than glossed through *post hoc* explanatory mechanisms. The interface must mirror the system's thinking. A rhetorical network is a coherent explanation for the claims made by the instantiated argumentation. The RST relations used are the same as those used to affirm coherence in naturally occurring texts. It should therefore follow that a rhetorical structure based on the argument ontology would possess a structure approximating that of a naturally occurring text. The explanatory power results not merely from the use of a rhetorical model, but through use of both inferential and synthetic argumentative modalities.

Through structured use of synthetic and inferential argumentation it is possible to assemble an explanatory discourse with rhetorical clarity. The example used is a simple one with limited argumentative depth. It also assumes that the segments used may be ontologically normalized. The basic idea here is that the network is the explanation.

Additional Areas of Research

Additional research areas include argumentative strategies and knowledge discovery. Fundamental argumentative strategies include *associative*, *inventive* and *inventive-dissociative*. *Associative strategies* play by the rules; they are applications of *modus ponens*, *modus tollens*, and other logical forms (Reed & Long 1998). Facts arrive and are assessed on the basis of argumentative structures. Reasoning proceeds by using warrants in establishing or denying connections between grounds or claims. Generative strategies invoke a warrant to affirm a connection between ground and claim. Destructive forms deny the claim to repudiate the grounds.

Inventive strategies do not deal with grounds or claims, but rather they challenge warrants, either by disputing the warrant directly or by challenging it with a new, competing warrant. Inventive strategies use arguments whose grounds and claims are themselves warrants. *Inventive-dissociative* strategies, mentioned earlier, tend to be

disruptive, as they introduce unfamiliar ontological structures into the discourse. Both inventive and inventive-dissociative strategies are based on Perelman and Olbrechts-Tyteca's concept of dissociative argumentation.

In an ontologically normalized knowledge base, discovery is possible through logical and analogical means. Logical discovery involves replacement operations such as commutation, association, distribution, and transposition. In analogical reasoning, claims are based not on an argument but on a situation deemed similar to the claim. What distinguishes analogical reasoning is the absence of any warrant. What gives the argument its persuasive power is the structural and ontological similarity between the analogues. If the similarity can be discovered by means of unification, it may also be possible to construct a warrant. By this means it may be possible to construct a new argument. In view of the rhetorical nature of the argumentative network, the resulting discoveries may be readily accessible for hum utilization.

Conclusion

Argumentative Reasoning Theory is a technology for knowledge representation. It defines a high fidelity model for support of intelligent collaboration between humans and computers. However, a model for intelligent human-computer collaboration is not merely a computer model. Several researchers have noted that successful collaboration requires establishment and sharing of common ground (Clark and Brennan 1991; Brock and Trafton 1999). Establishing common ground makes demands, not only of computers, but also of the people who use them. Common ground is achieved not solely by designing better technology. We also educate ourselves to accommodate to the strengths and weaknesses of technology. ART prescribes a common ground that would enable humans and computers to reason, create, learn, and discover together. The model employed is human on both sides of the equation, and in this way ART endeavors to offer an accommodation.

References

Brock. D and Trafton, J. G. 1999. Cognitive representation of common ground in user interfaces. In Kay, J. editor, *User Modeling: Proceedings of the Seventh International Conference, UM99* New York: Springer, 287-289.

Clark, H. H. and Brennan, S. E. 1991. Grounding in communication. In Resnick LB, Levine JM, Teasley SD editors, *Perspectives on Socially Shared Cognition Washington.* DC: APA Books, 127-149.

Fox, J. and Das, S. 2000. *Artificial intelligence in hazardous applications.* Menlo Park, CA: AAAI Press.

Grasso, F. 2002. Towards a framework for rhetorical argumentation. In Bos J, Foster M.E., Matheson C. editors, *6th Workshop on the Semantics and Pragmatics of Dialogue.* Edinburgh, UK, 53-60.

Juthe, A. 2005. Argument by Analogy. *Argumentation* 19: 1-27.

Mann, W. C. and Thompson, SA. 1988. Rhetorical structure theory: Towards a functional theory of text organization. *Text* 8:243-281.

Moulin, B., Irandoust, H., Bélanger, M., and Desbordes, G. 2002. Explanation and argumentation capabilities: Towards the creation of more persuasive agents. *Artificial Intelligence Review* 17: 169-222.

Perelman, C. and Olbrechts-Tyteca, L. 1969. *The new rhetoric: A treatise on argumentation.* Notre Dame, IA: University of Notre Dame.

Pollock, J. L. 1995. *Cognitive carpentry: A blueprint for how to build a person.* Cambridge, MA: MIT Press.

Reed, C and Long, D. 1998. Generating the structure of argument. *Proceedings of the 36th conference on Association for Computational Linguistics*. East Stroudsburg, PA: Association for Computational Linguistics. 1091-1097.

Sartor, G. 1993. A simple computational model for nonmonotonic and adversarial legal reasoning. *ICAIL '93: Proceedings of the 4th International Conference on Artificial Intelligence and Law.* New York: ACM Press. 192-201.

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

Wærn, Y and Ramberg, R. 2004. Distributed knowledge by explanation networks. *HICSS '04: Proceedings of the Proceedings of the 37th Annual Hawaii International Conference on System Sciences (HICSS'04) - Track 5.* Washington, DC: IEEE Computer Society.

Ye, L. R. 1995. The value of explanation in expert systems for auditing: An experimental investigation. *Expert Systems with Applications* 9: 543-556.

Ye, L. R. and Johnson, P. E. 1995. The impact of explanation facilities on user acceptance of expert system's advice. *MIS Quarterly* 19: 157-172.