AI Magazine Volume 16 Number 4 (1995) (© AAAI)

The Workshop on Computational Dialectics

R. P. Loui

■ The 1994 Workshop on Computational Dialectics was held during the 1994 National Conference on AI. At issue were the following ideas: (1) dialectic has something to do with computation, (2) AI has something new to contribute to the understanding of dialectic, and (3) dialectic approaches to long-standing AI problems permit new progress. This article outlines the significant research presented at the conference.

Tom Gordon (GMD) had the idea for an American Association for Artificial Intelligence (AAAI) workshop on computational dialectics, and the intriguing phrase is his. It is to his credit that the idea has earned a further hearing in a German workshop that he is currently organizing with colleagues on the continent. For the workshop in Seattle, Washington, at AAAI-94, our committee included Johanna Moore (Pitt) and Katia Sycara (Carnegie Mellon University). It is to their credit that they agreed to be on the committee, perceiving the fundamental nature of the workshop's topic and the broadness of its implications in time.

Senior colleagues Hajime Yoshino (Meijigakuin University) and Layman Allen (University of Michigan) were able to attend: Both have records of interdisciplinary leadership, where law and formal systems meet, and their input is always valued. Trevor Bench-Capon from England (Imperial College), Daniel Poulin from Quebec (Université de Montréal), and perhaps half of the other attendees made the trip to Seattle just for the workshop, although other scheduled speakers, notably Jaap Hage from The Netherlands (University of Limburg) and Arthur Merin from Germany (University of Stuttgart), were unable to attend. Gilad Zlotkin (Hebrew University), Sandra Carberry (University of Delaware), Gerhard Brewka (University of Vienna, Austria), Jon Doyle (Massachusetts Institute of Technology), Anne Gardner, L. Thorne McCarty (Rutgers University), and Kevin Ashley (Pitt) also contributed their intellectual weight to the proceedings.

At issue were the following ideas: (1) dialectic has something to do with computation, (2) AI has something new to contribute to the understanding of dialectic, and (3) dialectic approaches to long-standing AI problems permit new progress.

Dialectic is at the core of models of rational inquiry that are honest about procedural rationality. It is no accident that Herbert Simon and Nicholas Rescher, collaborating in the early 1960s, would shortly thereafter discover the idea of procedural rationality (Simon) and a compact logical description of dialectic (Rescher). Why should there be debate instead of just unpacking the declared logical entailments of nonmonotonic or defeasible reasons? The answer has to be that there is search, and the answer makes no sense unless search is limited. Simon and Rescher both knew this answer, but it has taken some time to put the full measure of their ideas together and make explicit pronouncements.

Dialectic … is basically the idea that rational inquiry is best achieved through largely adversarial discourse.

There were two notable forerunners of AI workshops that explicitly contemplated argument. The last was run successfully by Sergio Alvarado, who was able to interest a much larger number of researchers in natural language processing in the 1991 AAAI Spring Symposium Series. In contrast, the workshop in Seattle had a clear bias toward AI and law and knowledge representation. The population shift is natural: In the intervening years, formal models of argument have grown precise. They are much more precise than at any time in the few thousand years that Western intellectuals have studied argument as a logicolinguistic phenomenon. The AI and law community has felt the impact of the developments and has decided for some time to take the lead toward further developments. The organizers of the last Nonmonotonic Reasoning Workshop (Brewka) and the next Knowledge Representation and Reasoning Conference (Doyle) were present, and their presence is evidence that AI and law will be followed, if not joined, in its pursuits.

Dialectic has had a tough time being respectable in this century. In its more logically sober guises, it is basically the idea that rational inquiry is best achieved through largely adversarial discourse. Dialectic is an old idea that simply will not disappear. It is the idea of structured linguistic interactions proceeding according to a protocol. The term *computational dialectics* was meant to describe an area of activity in AI that considers the language and protocol of systems that mediate the flow of messages between agents constructing judgment, agreement, or other social choice to recognize or achieve an outcome in a fair and effective way.

Significant papers were presented by Brewka and McCarty: Brewka gave a formal reconstruction of Rescher's formal theory of disputation within the framework of default reasoning. If Rescher is indeed the high mark in the history of understanding dialectic, which now seems incontrovertible, then Brewka represents the first

direct step forward in nearly two decades. Of course, these two decades were active years for the nonmonotonic reasoning community that makes Brewka's work possible. Brewka uses the precision of his revised default framework to answer questions that Rescher did not or could not answer. I disagreed with Doyle over the long-term importance of precisely answering exactly these questions; I felt that variations of the model remain to be explored at a higher level of abstraction rather than to obsess over one variation's details.

Still, a full literature search of citations of Rescher's 1977 monograph, *Dialectics*, reveals no useful formal extension or clarification of the logical system prior to Brewka. The Argentine constructive logician Gino Roetti is the only person who has told me he read the book in 1977, saying he knew then that is was an important little book but that it was perhaps ahead of its time. Was an understanding of it really impossible until AI had had its nonmonotonic reasoning excursions?

McCarty's paper was as electrifying as it was an object of impassioned disagreement. Although primarily historical and not essentially concerned with the logic of dialectic, it was recognized at once for its importance. McCarty noted that it was the greatest Anglo-American philosopher of law, H. L. A. Hart, who introduced the logicolinguistic world to the term *defeasible* (Hart imported it from English contract law). These origins are perhaps of greater interest to the nonmonotonic reasoning community because this community has a stake in a better understanding of defeasible reasons; still, it is clear that dialectic presupposes defeasible reasons on which to construct arguments that can stand in opposition. McCarty raised the question of why Hart never used his term, *defeasiblity,* in his work again. There was superficial evidence that Hart even disavowed his early paper that introduced the idea of a defeasible concept, and the modern lore of jurisprudence holds that he did disown the paper and its ideas.

For Hart, the early idea of a defeasible concept probably evolved into his famous idea of an open-textured term to be defined perhaps using defeasible rules (which can be traced to the philosopher Ludwig Wittgenstein by way of philosopher Friedrich Waismann). However, McCarty took the opportunity to assail AI's work on defeasible reasoning. He said that defeasible rules were trivial devices that could be eliminated in favor of two kinds of negation. As devices for rule-based programming, they provide no advance over intuitionistic logic programming (says McCarty). They are trivial, that is, when compared to the defeasibility of open-textured concepts, the logic of which remains unanalyzed (says McCarty, downplaying the work of case-based reasoning researchers, especially his colleagues Edwina Rissland, Ashley, and David Skalak). Ironically, the AI and law luminary, McCarty, like the deontic legal logician Carlos Alchourron, chooses to attack a legal invention (defeasible reasoning) on technical aesthetics.

The issue of open texture was made more pointed by a paper presented by Pierre St-Vincent. He and Poulin gave a nondialectic, nondefeasible treatment of open-textured concepts. These Canadians wanted to analyze open texture as a normed or fuzzy concept, in the manner that might be appropriate at a meeting about uncertainty in AI. A role for dialectic remained in St-Vincent's approach: Dialectic provided for an exchange of messages between advocates interested in the contrary aims of enlarging the scope of a concept and restricting it. This paper would be an excellent place for fuzzy logicians to approach the issues raised by this workshop.

These papers, instead of pushing models of dialectic forward, involved debating the current understanding of defeasible reasoning and dialectic. The paper of Violetta Cavalli-Sforza and Dan Suthers and that of Kathy Freeman were counterpoised.

Cavalli-Sforza and Suthers reported on BELVEDERE, "an environment for practicing scientific argumentation." BELVEDERE is a tutoring system for high school and junior high school science classes. Graphic support is provided for diagramming arguments interactively. Like much of the work presented at the 1991 symposium, it extends the idea of Toulmin diagrams for arguments into the world of graphic user interfaces. It is ambitious work in part because all the research that builds on the GIBIS (Conklin) and AQUANET (Marshall et al.) experiences points to the average user's unwillingness to frame arguments formally. In addition, it insists that the arguments to be diagrammed be scientific arguments. Surely, scientific arguments have their own special logic. Cavalli-Sforza has for a while been interested in Toulmin's own attempts to apply his work on argument to specialized forms of reasoning, such as legal, decision-theoretic, or scientific. BELVEDERE is significant because it made good of Toulmin, with possible ramifications in the interactions of users with scientific databases; it is perhaps more significant in this respect than it is successful as a tool for teaching young scientists.

Freeman gave a synopsis of her dissertation work, which also included a graphic system for diagramming arguments based on Toulmin. Her system provides an automated reasoner for additional support of the user. Although her representational scheme seeks to mix dialectic ideas with measures of plausibility, the audience found nothing objectionable in her work. This reaction shows me that the research program that seeks to import graphic user interface versions of Toulmin's ideas to AI and computer-supported collaborative work (CSCW), a paradigm present in 1991 and mentioned in the recent Association of Computing Machinery computing-survey article on CSCW, is on strong footing.

Bench-Capon and P. H. Leng's paper was self-explanatory: developing heuristics for the argument-based explanation of negation in logic programs. The authors extend Bench-Capon's work presented at a series of expert system conferences: "Interacting with Knowledge-Based Systems through Dialogue Games (with P. E. S. Dunne and P. H. Leng, 1991), "Using Toulmin's Argument Schema

to Explain Logic Programs" (with D. Lowes and A. M. McEnery, 1991), "A Dialogue Game for Dialectical Interaction with Expert Systems" (with P. E. Dunne and P. H. Leng, 1992), and "Argument-Based Explanation of the British Nationality Act as a Logic Program" (with F. P. Coenen and P. Orton, 1993). The premise in all this work is that negation as failure is inadequate as a report to a user. In its stead, a dialogue game is played in which the automated reasoner uses its recently completed search to defeat a player who seeks to establish the contentious proposition (against the advice of the program). The user's stubbornness as a player in such a game defines the amount of information that is required to explain the result: the more stubborn the use, the longer the dialectic game, the larger the explanation generated. Understanding the interaction as dialectic, as a two-player adversarial game of responses, is novel. It is what computational dialectics brings to AI's foundations, with ramifications in the design of systems.

I believe that Allen's paper, Gordon and Brewka's paper, and my own could also be regarded as novel. Allen gave the rules of an A-Hohfeldian game, which some might properly regard as a *Wff'n'Proof* (a 1960s board game that taught logic) for logical statements using the turn-of-the-century, heretofore underappreciated, largely deontic ontology of the legal scholar Wesley Hohfeld. The work was jointly undertaken with Charles Saxon. The Allen-Saxon revision of Hohfeld's ontology first saw good use in an expert system reported at the 1991 First International Workshop on Deontic Logic. In our Seattle workshop, the work took a turn toward formal games, renewing Allen's interest in producing the rules of adversarial logic games. Allen hinted that rules for a fully disputational discourse game, that is, his ideas for a model of dialectic disputation, were in the works.

Gordon and Brewka gave a paper about decision making. The preference ordering required for evaluating trade-offs in decision making was assimilated into the ordering used to represent which arguments are preferred over their opposing arguments. Some people will think their analysis is backward, but this is the novelty. Past authors have imported the axioms of utility into the analysis of multiple nonmonotonic extensions (M. Wellman and J. Doyle, "Impediments to Universal Preference-Based Default Theories, *AI Journal* 49(1): 97–128). The idea here is to use the machinery of defeat among arguments to replace real-valued utility when decision making requires making trade-offs. Instead of saying that money and time can be mapped to utilities, say that an argument for a decision based on consequences of one kind (one attribute) defeats an argument for a different decision based on consequences of some other kind (other attribute). It is a qualitative approach to multiattribute utility.

The idea seems initially sound, although it is perhaps pushed too far. Preferences among arguments are sometimes based on superficial syntactic features, such as specificity among defeasible reasons. These preferences might be used purely for notational convenience. If so, then the availability of only one kind of preference might be representationally impoverished rather than merely clever.

My paper, "Argument and Arbitration Games," sought to include an argument game within the various games of negotiation. The approach of game theory to negotiation is to characterize it as a series of proposals. Sometimes settlement is based on reason, although as Sycara notes, sometimes the reason is merely that one person is in a better bargaining position than the other. It is unfortunate that Arthur Merin was unable to give his paper along with mine because the two seemed to be in line with each other. The work has just begun, but I am already perched on the limb that holds that future economists' models of negotiation will integrate discourse and reason. Cheap talk and precedent are simply too important as parts of what can be considered negotiation.

Zlotkin gave a synopsis of his work and his reaction to the thrust of the workshop. He and Stan Rosenschein had just produced the nifty monograph *Rules of Encounter* (The MIT Press, 1994), in which the economic approach to achieving agreement is used to derive interesting results in distributed AI. Compared to their methodology, the emerging methodology of computational dialectics seems troubling. Rosenschein and Zlotkin prove interesting theorems about games; computational dialecticians give rules for interesting games.

Understanding the interaction as dialectic, as a two-player adversarial game of responses, is novel. It is what computational dialectics brings to AI's foundations, with ramifications in the design of systems.

Providing rules for a game without justifying or analyzing these rules might be the right thing to do and might be no different (in the abstract) from what people do when they design AI programs. Describing such discourse games leads to neither theorem nor consensus. It just describes a regimen for the interaction of a society of minds. Clearly, this regimen is important, but there must be some way to evaluate the merit of the resulting models. Computational dialecticians can defend their results and develop the kinds of taxonomy that logicians have for their wares. They can also inspire better programs, as seen in the pre-

sentations of Cavalli-Sforza and Suthers and Bench-Capon and Leng. More work needs to be done.

We had our share of students completing dissertations: George Ferguson (under James Allen at Rochester), Jennifer Chu-Carroll (under Carberry at Delaware), Vincent Aleven (under Ashley at Pitt), and Jeremy Wertheimer (with John Mallery at MIT). Although each has, in the past, written papers and programs that were centrally relevant to the theme of the workshop, all remained surprisingly quiet. This problem is a major failing of a workshop. Although the topic and its discussion heavily favored law and logic interdisciplinarians, there is no future for a field that does not promote the development of its future scholars.

The biggest disappointment was that Gerard Vreeswijk could not attend because his dissertation work neatly defines the area. His work seeks explicitly to model dialectic as a discourse game in logic. He is now taking the same paradigm of computational dialectics and addressing distributed AI problems of coordination within it. I see Vreeswijk's remaining on the other side of the Atlantic Ocean as a metaphor. Much of the best work on this topic will likely remain on the continent, where the temperament is more intellectual, the logical traditions are more varied, and the relevance of paradigm to application is better welcomed. Like AI and law, the conditions for leadership exist outside the United States. We were lucky in Seattle that the American conference was able to take the lead for a year.

R. P. Loui is associate professor of computer science, affiliate associate professor of legal studies, and research affiliate of the Center for Semantic Control and Optimization at Washington University in St. Louis, Missouri.