CHAPTER

24

The Impact of Argumentation on Artificial Intelligence

Douglas Walton
University of Winnipeg
David M. Godden
University of Windsor

The research of the Amsterdam School has spread outward across the discipline of argumentation studies like a new day, awakening us to new vistas, casting light on new opportunities, and offering a fresh look at our familiar surroundings. When it first appeared, the pragma-dialectical approach challenged so many existing assumptions that is seemed almost radical, and entirely disrupted the established view. Yet over the years this approach has proved so remarkably effective that many of its central tenets have begun to be widely recognized and accepted. These tenets are even becoming a part of science, as they are increasingly adopted into the standard model of argument used in computing. Along with Rob Grootendorst, Frans van Eemeren was the founding father of the Amsterdam School, and of the pragma-dialectical approach to the study of argumentation. This new approach found its inspiration in the critical rationalism of Popper (1972, 1974), Barth and Krabbe's (1982) theory of formal dialectic, and the speech act theory of Austin (1962), Searle (1969), and Grice (1975) (van Eemeren & Grootendorst, 2004, p. 51).

Argumentation, as a growing interdisciplinary field of research, was conducted mainly in logic, philosophy, and communication studies in the beginning. It has now branched and become truly interdisciplinary as it has affected more and more fields, like cognitive science, where models of rational thinking are an essential part of the research program. At some point, argumentation methods and findings began to be imported into computing, especially in the area called artificial intelligence, or AI. Since that time, other researchers in argumentation began to use tools developed in AI.

In this chapter, we explore the development and importance of this connection between argumentation and artificial intelligence. Specifically, we show that the influence of argumentation on AI has occurred within a framework that is consistent with the basic approach of Pragma-Dialectics. While the pragma-dialectical approach is typically conceived

of as applying primarily to argumentation occurring between human agents, we show that the basic features of this approach can consistently be applied in a virtual context, whereby the goal-directed activities of, and exchanges of information between, artificial agents are regulated by procedural rules.

PRAGMA-DIALECTICS: BASIC FEATURES

Most readers will already be familiar with the pragma-dialectical (PD) approach to the study of argument. In this section we briefly review some of its characteristic features for the purpose of showing how they have been adopted in AI.

The observation that arguments are primarily communicative activities occurring between rational agents yields the central PD insight that the subject matter of argumentation can be conceived of as a complex speech act which is both constituted and regulated by pragmatic rules. The pragma-dialectical methodology unifies the normative and descriptive approach to the study of argumentation (van Eemeren *et al*. 1993). In contrast to *audience-oriented* rhetorical models, the pragma-dialectical model is *resolution-oriented* (van Eemeren, 1994, pp. 6-7) and reflects a *critical-rationalist* approach towards its subject matter (van Eemeren, 1994, p. 4). In contrast with epistemic models, pragmatic rules such as Grice's Cooperative Principle (1975, p. 45) and the conversational maxims of Quality, Quantity, Relation, and Manner provide both the format and the foundation of the norms governing argumentation.

Central to the pragma-dialectical conception is that argumentation is a verbal, social, and rational activity. Van Eemeren and Grootendorst (2004, p. 52ff) described the pragma-dialectical theory as being founded on the four meta-theoretical principles of externalization, functionalization, socialization, and dialectification. Externalization stresses the external and the explicit over the internal and the implicit. It involves focusing on the public commitments of arguers arising from their linguistic activity, rather than their unexpressed, private beliefs. Functionalization stresses the goal-oriented and procedural aspects of argumentation. It involves adopting a process-based, or functional, view of argument over a product based, or structural, view. The linguistic activity of argumentation is treated pragmatically as a regulated sequence of purposive speech acts. Socialization highlights a picture of argumentation which is interactional and has multiple agents, rather than one which is individual and has a passive-audience. It involves treating argumentation as a dialogue between two or more parties with distinguishable and opposing positions, each of whom actively participate in the argumentation process. Dialectification involves treating dialogue moves in a normative context, as attempts to resolve a difference of opinion in accordance with critical norms of reasonableness. This allows evaluative principles to be stated as procedural norms regulating the activity of argumentation, rather than as standards against which a product is later measured. Together these principles inform both the methodology of Pragma-Dialectics, and its ideal model of argumentation.

A central motivation of the Amsterdam School seems to be the rationalization of our

social institutions and the development of a more reasonable population. For example, van Eemeren (1995, p. 145) has noted the importance of argumentation to democracy, by ensuring that it is both a participatory and rational process. Indeed, the achievement of these noble goals has been remarkably successful, especially in Holland where the pragma-dialectical tools and perspective have received recognition and acceptance even beyond the halls of academia in fields like law and public policy. Less studied, though, is the application of the PD approach in virtual, or nonhuman contexts. Yet, if we are right, the usefulness of PD in fields such as AI could have an even greater impact.

While the full impact of PD on computing has yet to be made apparent, in this chapter we briefly describe some of the recent developments in AI that have led to its initial collaboration with argumentation. In studying this connection, we note that many of the basic argumentative principles adopted by AI are well fitted to the pragma-dialectical approach. Specifically, we argue that the approach to multi-agent systems and computational dialectics currently prevalent in AI is consistent with the basic meta-theoretical principles that inform the pragma-dialectical approach to argumentation. In doing so, we observe some of the existing developments in this collaborative field of research that could help to lay the groundwork for future advances.

AI AND ARGUMENTATION: A BRIEF INTRODUCTION

Most people know what artificial intelligence is, and why it is an important branch of computing. When research in AI first began, the problem seemed to be that in order to build useful, effective AI systems we had to try to understand how to automate a reasoning process that approximates (if not replicates) human thinking. This project turned out to be a lot harder than anyone thought. While it would appear that modern computers have more raw computational power than the human brain, this does not seem to be the limiting factor. Research has shown that common sense reasoning of the kind used to carry out the many mundane tasks we perform every day is based on a lot of implicit knowledge about how things work in our social and physical world. Supplying a computer with this seemingly ad-hoc list of common knowledge, gained through a lifetime of experience and learning, has proved a Sisyphean task. Another problem was that deductive formal logic, of the kind emphasized in science, philosophy, and especially mathematics, didn't seem to help as much as many thought it should. Although this led some AI researchers to abandon the project of modeling reasoning, it led others to look beyond the bounds of deductive logic. Deductive logic is based on quantifiers like 'for all x' that do not admit of exceptions, while the kind of reasoning needed for AI is defeasible, meaning it is based on rules that are subject to exceptions, producing inferences that sometimes default.

Argumentation was friendly to this notion of reasoning subject to default, based as it was on the notion that argumentation needs to be analyzed and evaluated in the framework of a critical discussion with two sides. The idea is that the purpose of argumentation

is to resolve a conflict of opinions, not exclusively by one side's producing arguments that are defeasible, but also by the other side's producing doubts, objections, and critical questions that probe into the weak points in the argument. This dialectical approach was helpful for dealing with inconclusive but nevertheless reasonable arguments like argument from expert opinion. This form of argument, although far from unconditionally trustworthy and even traditionally classified as a fallacy, can often be reasonable as a way of reaching a decision, inferring a tentative conclusion, or justifying a standpoint under conditions of uncertainty. After all, we can't all be experts on everything, and, in order to be reasonable, we sometimes have to depend on the opinion of an expert, even though we need to be critical of these arguments as well. For example, expert testimony is rightly treated as an important form of evidence in law, even though the acceptability of expert testimony involves its being subjected to cross-examination at trial.

FROM ARGUMENTATION TO AI

It is somewhat unusual that interdisciplinary research conducted mainly in subjects in the humanities and social sciences would become useful in a fundamental way in an important area of research in the natural sciences. But that is what has happened with argumentation. It started as an interdisciplinary effort converging around the meetings of ISSA (the International Society for the Study of Argumentation), the Summer Conferences on Argumentation in Alta, Utah (hosted by the National Communication Association and the American Forensic Association) and an informal logic group mainly centered in Canada. Many of the leading exponents were professors of philosophy or speech communication. In Europe, the research clustered around the impetus of the Amsterdam School, who had built up a pragma-dialectical theory of argumentation attracting a small but dedicated group of adherents. In 1987-88, a research group on fallacies worked together for 1 year at the Netherlands Institute for Advanced Studies in the Humanities and Social Sciences (NIAS).

As the ISSA Conference met every 4 years, and spin-off conferences in North America like the OSSA Conference built up momentum, a curious thing happened. A small group of researchers in computing began to pay attention to their results. The first early meeting was the June, 1996 International Conference on Formal and Applied Practical Reasoning held in Bonn, Germany. We share the opinion of Verheij that this conference was of unique importance. In his report of this conference, Verheij (1996) described its purpose, which made an explicit statement of the goals and advantages of cooperation between the two groups:

> One of the problems of modern scientific research is that it almost inevitably requires a high level of specialization. The advantage is that the resulting ivory towers, in which most individual researchers spend their scientific lives, can more efficiently be built higher; the disadvantage is that these towers do not simply add up to an integrated scientific village .

. . Research on argumentation can have rather different aims, for instance:

- To describe and evaluate actual human argumentation by means of empirical investigation (e.g., in cognitive science or psychology).
- To apply an argumentation model in order to build intelligent computers and programs (e.g., in computer science and artificial intelligence).
- To investigate and enhance our conceptualizations of argumentation in order to better understand its nature (e.g., in philosophical and mathematical logic). (Verheij, 1996)

More recently, the impetus toward collaboration between the groups culminated in the Bonskeid House Conference which took place in Pitlochry (Scotland) in 2000. At this conference, a selected group of researchers, some from computing and some from argumentation, met together in an isolated castle in the highlands, to try to find out how each field could benefit from the findings and methods of the other. Its stated objectives were the following:

1. Bringing together a small number of researchers from several different fields to encourage the exchange of new ideas and the establishment of new collaboration.
2. Producing a handbook authored jointly by all participants which will serve to disseminate to a wide audience the defining problems, issues, and avenues for future work.
3. Introducing a small number of high quality research students to collaborative ventures between computational sciences and argumentation, and to encourage them to pursue their own research interests with an awareness of the ties with areas of mathematics, philosophy and AI. (Norman, 2000)

The *modus operandi* of the conference was to divide up the participants into five small groups, working on the following topics: argument and computational societies, argument and practical reasoning, argument and legal reasoning, argument and computational linguistics, and computational models of argument. Each group had the rather formidable job of writing a chapter of a book over the course of the week. As a result of this feverish effort, the book *Argumentation Machines: New Frontiers in Argument and Computation* (Reed & Norman, 2004) was eventually produced.

COMPUTATIONAL DIALECTICS

As initial efforts to form partnerships and set out a common research agenda gained momentum, a new field of computing arose called computational dialectics. The need for a field like computational dialectics arises from the situation created by multi-agent systems, or distributed systems, in computing. Basically, multi-agent systems are composed of a number of autonomous agents interacting in a computerized environment in which there is no centralized control. Agents can be artificial (like computer programs or robots) or natural (like humans), and are characterized by certain features. Agents are purposive

and they are intelligent. That is, they can act and interact in the system, and have goals which can be individual or communal. Also, they can store limited amounts of information and 'sense' or receive information from their environment. Typically, agents do not initially have sufficient 'knowledge' to fulfil their goals, and so they must find ways of interacting with one another to do so by collecting the required knowledge.

Importantly, multi-agent systems have become ubiquitous in the computing world, and characterize many online environments such as expert systems, e-commerce, and legal-support systems. The development of multi-agent systems for communication on the web requires a framework in which intelligent agents can engage in various kinds of conversational interactions through speech acts. These include asking questions to get information from other agents, assessing the worth of that information as a basis for arriving at an intelligent decision, and generally, reasoning together to solve a problem or resolve a disagreement. Faced with the situation created by multi-agent systems, designers and engineers began work on computational dialectics as a means of regulating the operations of distributed systems and the interactions of the agents of which these systems are composed.

Thomas Gordon (1996) describes the field of computational dialectics in the following way:

> The subject matter of this field [of computational dialectics] is the design and implementation of computer systems which mediate and regulate the flow of messages between agents in distributed systems, so as to facilitate the recognition and achievement of common goals in a rational, effective and fair way.

What is interesting about this new field of computing is that it is intrinsically dialectical, and is significantly informed by work in argumentation.

According to Lodder (2000, p. 255), the term *computational dialectics* first appeared when Ron Loui and Tom Gordon organized an American Association for Artificial Intelligence (AAAI) workshop with Johanna Moore and Katya Sycara under the name Computational Dialectics in Seattle in 1994. The call for papers for the workshop (Loui & Gordon, 1994) described the field of computational dialectics as the study of structured dialogues used in multi-agent communication systems in which agents reach agreement through rational interaction in order to achieve common goals in a fair and effective way. Such dialogues contain a blend of adversarial argumentation, so that each agent has an individual goal, and is an advocate, but at the same time the procedure is collaborative and works only because the agents also share a common goal. Thus the workshop invitation used the expression 'communal standards' when describing such computational models of deliberation, negotiation, and discussion.

The centrality of the use of structured dialogues as a means of regulating communication and of rationally achieving common goals in distributed systems indicates exactly how much the very idea of computational dialectics draws upon the dialectical conception of argumentation, and on the theoretical resources it has to offer. Consider the

description of the field as presented for the 1994 AAAI Workshop on computational dialectics, which is worth quoting at some length.

> Dialectic is an idea that simply will not disappear. It is the idea of structured linguistic interactions proceeding according to a largely adversarial protocol. Beginning with the ancients, dialectic appears to many to be synonymous with rationality. Today, computation informs the study and use of such structured dialogues. Substantial contributions are now possible from artificial intelligence researchers.
>
> The term "Computational Dialectics" is meant to describe an area of activity in AI, which 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.
>
> The study of communal standards for acquiring knowledge and making decisions has always been interesting as a basis for computational models of deliberation. The study of argument and negotiation in naturally occurring dialogues has been the focus of work in language processing and explanation generation. The study of analogy and case-based reasoning has produced dialectical models that have been successfully applied in the domain of conflict resolution, negotiation and legal argument. The implementation of non-monotonic reasoning systems and the semantics of logic programming has also converged on dialectic. Philosophers' formalization of defeasible reasoning has produced new understanding of why dialectic is not merely roundabout proof. Researchers of HCI [human computer interaction] and CSCW [computer supported cooperative work] have had occasion to study how the interplay of argument, counterargument, and rebuttal affects design, clarifies presentation, and improves interaction. (Loui & Gordon, 1994)

In addition to a sound appreciation of the dialectic approach as something more than 'roundabout proof' but as essentially involving rebuttal and counterargument, several other points are worthy of note in the foregoing account. First, it is interesting that one of the initial problems facing the field of AI has been overcome through the use of defeasible rather than deductive logics in these dialectical systems. Also interesting is the recognition of various kinds of dialogues such as deliberation and negotiation which become contexts of argumentation. But what about the speech acts of convincing and persuading which are so central to the pragma-dialectical framework?

The Intelligent Systems Group at the University of Utrecht explains how computational interactions on the web involve speech acts of convincing and persuading: "Many decisions are not made individually, but in interaction with others (for example in discussions, meetings, negotiations, and legal procedure). In such environments it is important to bring different opinions together and convince or persuade each other, instead of merely exchanging information" (University of Utrecht, 2005). Here one can see why argumentation has found such wide acceptance in computing, and why it fits there so naturally.

COMPUTATIONAL DIALECTICS AND PRAGMA-DIALECTICS: CONNECTIONS

While the foregoing has only been a brief survey of the historical developments and present status of research occurring at the intersection of argumentation and AI, it provides sufficient background to demonstrate the connections between this line of research and the pragma-dialectical approach. Having introduced some of the characteristic features of computational dialectics, it remains to show how these are in principle consistent with the fundamental principles of Pragma-Dialectics. It should be clear that the basic similarities of approach are overwhelming, or at least that they are sufficient to warrant a closer look. In this section, we explore the consistency of the four basic meta-theoretical principles of Pragma-Dialectics—socialization, functionalization, externalization, and dialectification—with the approach adopted in computational dialectics.

Socialization. Perhaps the most obvious point of connection between computational dialectics and the pragma-dialectical approach to argumentation is socialization. Socialization, in Pragma-Dialectics, involves conceiving of argumentation as a communicational activity that occurs between two or more active participants. It leads directly to the conception of argument as a dialogue, and to a dialectical approach to its study. Such a perspective clearly applies to multi-agent systems in AI. The communicative interaction of a number of intelligent, purposive agents forms the very basis of multi-agent systems.

Functionalization. Closely related to socialization is the idea of functionalization. Central to the functionalization of argumentation is adopting a processes-based view of it, whereby the activity of argumentation is seen as the result of a goal-oriented, communicative action. Together with the idea of socialization, functionalization leads to the idea that we can have a procedural rather than a structural approach to the regulation of argumentation. Given the earlier description of multi-agent systems, its consistency with the pragma-dialectical tenet of functionalization should be fairly clear. As mentioned, agents themselves are treated as purposive entities who can act to attain their own individual goals as well as well as shared goals. Furthermore, the communicative activities that occur within multi-agent systems occur in the context of this goal-fulfilling process.

Externalization. The basic idea of externalization is that argumentation theory must start with what is explicit (i.e., those commitments that can be publically attributed to an arguer, and to which the arguer can be held publically accountable). This basic principle of working with commitments as opposed to beliefs is a point that is shared by most dialogue-based approaches to argumentation (e.g., van Eemeren *et al*. 1993, Walton & Krabbe, 1995; van Eemeren & Grootendorst, 2004), and is important not only from the point of view of the theorist, but from the perspective of the arguers also. Starting from publically-accessible commitments ensures that arguers 'have something to grab on to' when recognizing differences of opinion and in identifying the contents of opposing standpoints. It also ensures that there is a set of shared, or mutually agreed upon premises, which can be used as a common starting point by participants in argumentation. (Or at least, that arguers can check to see whether there is such a common starting point for further discussion.)

Importantly, this idea of commitment has contributed significantly to the prevailing conception of multi-agent systems in AI. Singh (1996, 1997) for instance, has suggested that multi-agent systems can usefully be conceived of as 'spheres of commitment.' And, the nature of commitments as they occur in AI bears a striking resemblance to the way commitments are conceived of in argumentation. Singh (1999) describes commitments as having the following properties: They are (a) attached to autonomous, rational agents, (b) inherently external and social, (c) revokable but binding upon those agents; and finally that they are (d) incurred and retracted through a process of social interaction which is (e) regulated by a set of procedural rules governing that interaction. Consequently, the principle of externalization is consistent with the typical operations of multi-agent systems in computing, and with computational dialectics.

Dailectification. Dialectification is the entry point for the normative dimension of Pragma-Dialectics, the essence of which involves using the ideal model of a critical discussion (van Eemeren & Grootendorst, 2004, chap. 3) in the practice and evaluation of argumentation. The normative status of this ideal model is justified by its problem validity (or problem-solving validity) and intersubjective validity (or conventional validity) (van Eemeren & Grootendorst, 2004, pp. 16-17, 57; cf. Barth & Krabbe, 1982, pp. 21-22). We show that similar standards are used to justify the models used in computational dialectics.

Intersubjective validity is an anthropological standard of reasonableness, and "has to do with the conformity between the [ideal] model's components and the values, standards, and objectives actual arguers find acceptable" (van Eemeren *et al*., 1993, p. 14). Such a standard sounds reminiscent of the "communal standards for acquiring knowledge and making decisions" (Loui & Gordon, 1994) used in computational dialectics. Recall also the basic similarity of situation envisioned by both Pragma-Dialectics and computational dialectics whereby there is no 'external judge' in the situation of argumentation to impose a standard or decision. Rather, the actual participants to the discussion must come to determine, and agree upon, those standards on their own.

Problem validity is an instrumental standard of how successfully the rules of the ideal, or regulative model bring about the goals of the discussion. To see how this standard is implemented in computational dialectics, consider Prakken's (2000) statement concerning the proper design of dialogue systems.

> In the present framework, the initial situation of a persuasion dialogue is a conflict of opinion between two rational agents about whether a certain claim is tenable, possibly on the basis of shared background knowledge. The goal of a persuasion dialogue is to resolve this conflict by rational verbal means. The dialogue systems [used to achieve this goal] should be designed such that they are likely to promote this goal.

Not only does Prakken adopt the standard of problem validity for dialogue systems, he also appears to have incorporated the basic fundamentals of a critical discussion into his framework. For example, his description of the initial situation closely matches the confrontation stage of a critical discussion.

A further consequence of the dialectification of argumentation is to analyze fallacious argumentation procedurally, as moves in a dialogue that impede the resolution of the dispute (van Eemeren & Grootendorst, 1987). This idea is easily reflected in computational dialectics, since the rules governing such systems are inherently procedural to begin with.

As of yet, and to the best of our knowledge, there has been no attempt to explicitly structure the regulatory procedures used in AI and computational dialectics into the four familiar stages of confrontation, opening, argumentation, and concluding (van Eemeren & Grootendorst, 1984, pp. 85-88), nor to implement the specific rules characteristic of the pragma-dialectical critical discussion (van Eemeren & Grootendorst, 1984, chap. 7; van Eemeren & Grootendorst, 2004, chap. 6). Yet, we see no reason, in principle, why such a normative structure could not be adapted to, and implemented in, distributed systems as a means of regulation and rationalization. Clearly, the project of doing so is well beyond the scope of this chapter.

What is important to note, though, is that similar systems already exist at various stages of development, and many of the crucial topics have already begun to be addressed. For example, Norman, Carbogim, Krabbe, and Walton (2004) have shown how many issues central to the dialectical approach to argumentation relate to multi-agent systems in computing. Hitchcock, McBurney, and Parsons (2001) have provided a set of procedural rules that could effectively regulate deliberation dialogues. Walton and Godden (forthcoming) have addressed the topic of embedding persuasion dialogues in negotiation dialogues. Bex, Prakken, Reed, and Walton (2003) and Reed and Walton (2005) have shown how argumentation schemes can be incorporated into such systems.

CONCLUSIONS

The impact of argumentation on AI has been considerable already, but shows great promise of increasing exponentially as new applications are found. Also notable has been the reverse process. New software tools for argument diagramming like *Aracuaria* (Reed & Rowe 2001) are revolutionizing the field of argumentation by providing automated resources that are useful for the analysis and evaluation of argumentation. Not only are these tools useful for helping to teach critical thinking to students at all educational levels. They are also proving to be useful in many other areas, for example, as applied to legal argumentation (Bench-Capon, 1997, 2002; Bench-Capon & Sartor, 2003; Bex & Prakken, 2004). Who would have thought in those early days when researchers from all over the world were getting together at the early conferences more than two decades ago that things would have gone in this direction?

Indeed, one of the key motivating concerns shared by all in those early moments of our discipline, that of antiformalism, seems to be being challenged by the applicability of argumentation in AI. The implementation of the tools of argumentation in AI and other computerized models requires that they be formalized to the degree that they can be mechanically executable. Yet, this new formalism reflects many of the insights produced

by its anti-formalist contributors. In the first place, current artificial models allow for defeasible argumentation by working with non-monotonic logics that better reflect the actual conditions under which everyday inferences and arguments are normally made. Furthermore, contemporary approaches do not model the structure of argument as a proof, or a sequence of theorems derived from a set of axioms by a set of rules. Rather, this new formalism has a more pragmatic orientation; what is modeled is a sequence of moves, or actions, by rational agents, and the rules of the system govern these moves. Moreover, in the case of multi-agent systems, these rules regulate the interactions of several agents with one another. As such, these models use a procedural approach to the regulation and codification of argumentation as opposed to an epistemic approach, and inherently allow for a dialectical conception of argumentation. So, not only has argumentation contributed to the development of AI, but this liaison can help argumentation in its own development and self-conception. It would seem that as greater progress is made in the study of argumentation, the more precisely we are able to articulate the principles and standards of the discipline, and the closer we come to being able to formalize those principles.

Still, and for all that, argumentation and computational dialectics remain outside the mainstream in fields like philosophy, law, and computing. In the mainstream the models of rational thinking that still dominate are deductive logic, inductive reasoning (the Bayesian model especially), and the cost-benefit type of rational decision making that has been so central in economics in the past. With only very few exceptions like the University of Amsterdam, argumentation is not taught in graduate schools. Instead, in philosophy, formal deductive logic is what is thought to be exclusively necessary for graduate students to master, while informal logic is by and large relegated to the poor stipends who must teach the mass introductory lectures on critical thinking.

There has been talk lately, however, about a new paradigm of rational thinking in the social sciences. In economics, bounded rationality is now widely accepted as the model of reasoning that should be applied to economic phenomena like consumer decision making (O'Driscoll & Rizzo, 1985). To cite another example, in both computing and the social sciences, the importance of the argument from ignorance, or lack of evidence argument, has been recognized. These developments suggest that the impact of argumentation has been considerable, not only on AI, but on many fields that use AI as a model that can be applied to all kinds of cases of human decision making, deliberation, and rational thinking under conditions of less than perfect knowledge.

ACKNOWLEDGMENTS

Research for this paper was made possible by separate research grants from the Social Science and Humanities Research Council of Canada held by each to the two authors.

REFERENCES

Austin, J. L. (1962). *How to do things with words*. Oxford, England: Clarendon Press.

Barth, E. M., & Krabbe, E. C. W. (1982). *From axiom to dialogue*. Berlin: de Gruyter.

Bench-Capon, T. (1997). Argument in artificial intelligence and law. *Artificial Intelligence and Law*, *5*, 249-261.

Bench-Capon, T. (2002). Agreeing to differ: Modelling persuasive dialogue between parties with different values. *Informal Logic*, *22*, 231-245.

Bench-Capon, T., & Sartor, G. (2003). A model of legal reasoning with cases incorporating theories and values. *Artificial Intelligence*, *150*, 97-143.

Bex, F., & Prakken, H. (2004). Reinterpreting arguments in dialogue: An application to evidential reasoning. In T.F. Gordon (Ed.), *Legal knowledge and information systems* (pp. 119-130). Amsterdam: IOS Press.

Bex, F., Prakken, H., Reed, C., & Walton, D. (2003). Towards a formal account of reasoning about evidence: Argumentation schemes and generalizations. *Artificial Intelligence and Law*, *11*, 125-165.

Eemeren, F. H. van. (1994). The study of argumentation as normative pragmatics. In F. H. van Eemeren, & R. Grootendorst (Eds.), *Studies in Pragma-Dialectics* (pp. 3-8). Amsterdam: Sic Sat.

Eemeren, F. H. van (1995). A world of difference: The rich state of argumentation theory. *Informal Logic*, *17*, 144-158.

Eemeren, F. H. van, & Grootendorst, R. (1984). *Speech acts in argumentative discussions: A theoretical model for the analysis of discussions directed towards solving conflicts of opinion*. Dordrecht: Foris.

Eemeren, F. H. van, & Grootendorst, R. (1987). Fallacies in a pragma-dialectical perspective. *Argumentation*, *1*, 283-301.

Eemeren, F. H. van, & Grootendorst, R. (2004). *A systematic theory of argumentation: The pragma-dialectical approach*. Cambridge, MA: Cambridge University Press.

Gordon, T. F. (1996). Computational dialectics. In P. Hoschka (Ed.), *Computers as assistants-A new generation of support systems* (pp. 186-203). Hillsdale, NJ: Lawrence Erlbaum Associates. Also can be retrieved (May 1, 2005) from www.tfgordon.de/publications/Gordon1996a.pdf

Grice, H. P. (1975). Logic and conversation. In P. Cole, & J. Morgan (Eds.), *Syntax and semantics. Vol. 3: Speech acts* (pp. 41-58). New York: Academic Press.

Lodder, A. R. (2000). Book review of Gordon (1995). *Artificial Intelligence and Law*, 8, 255-264.

Loui, R. P., & Gordon, T. F. (1994). *Call for papers for 1994 AAAI Seattle workshop on computational dialectics*. Retrieved May 1, 2005, from http://www.cs.wustl.edu/~loui/comectics.text

Norman, T. (2000). *Symposium on argument and computation* [unpublished conference report]. Retrieved May 1, 2005, from http://www.csd.abdn.ac.uk/~tnorman/sac/

Norman, T. J., Carbogim, D. V., Krabbe, E. C. W., & Walton, D. (2004). Argument and multi-agent systems. In C. Reed, & T. J. Norman (Eds.), *Argumentation machines: New frontiers in argument and computation* (pp. 15-54). Dordrecht: Kluwer.

O'Driscoll, G. P. Jr., & Rizzo, M. J. (1985). *The economics of time and ignorance*. London: Routledge.

Popper, K. R. (1972). *Objective knowledge: An evolutionary approach*. Oxford: Clarendon Press.

Popper, K. R. (1974). *Conjectures and refutations: The growth of scientific knowledge*. London: Routledge & Kegan Paul.

Prakken, H. (2000). On dialogue systems with speech acts, arguments, and counterarguments. In M. Ojeda-Aciego, I. P. de Guzman, G. Brewka, & L. Moniz Pereria (Eds.), *Proceedings of JELIA'2000, The 7th European workshop on logic for artificial intelligence, Springer lecture notes in AI* (pp. 224-238). Berlin: Springer Verlag.

Reed, C., & Norman, T. J. (Eds.). (2004). *Argumentation machines: New frontiers in argument and computation*. Dordrecht: Kluwer.

Reed, C., & Rowe, G. (2001). *Araucaria*. Software. Retrieved May 1, 2005, from http://www.copmputing.dundee.ac.uk/staff/creed/araucaria/

Reed, C., & Walton, D. (2005). Towards a formal and implemented model of argumentation schemes in agent communication. In I. Rahwan, P. Moraitis, & C. Reed (Eds.), *Argumenation in multi-agent systems* (pp. 19-30). Dordrecht: Srpinger.

Searle, J. R. (1969). *Speech acts: An essay in the philosophy of language*. Cambridge, England: Cambridge University Press.

Singh, M. P. (1996). Multiagent systems as spheres of commitment. *International conference on multiagent systems (ICMAS), Workshop on norms, obligations, and conventions.* Kyoto, Japan. Retrieved May 1, 2005, from http://www.csc.ncsu.edu/faculty/mpsingh/papers/mas/icmas-96-norms.pdf

Singh, M. P. (1997). Commitments among autonomous agents in information-rich environments. *8th European workshop on modelling autonomous agents in a multi-agent world*, Ronneby, Sweden. Retrieved May 1, 2005, from http://www.csc.ncsu.edu/faculty/mpsingh/papers/databases/maamaw-97.pdf

Singh, M. P. (1999). An ontology for commitments in multiagent systems: Toward a unification of normative concepts. *Artificial Intelligence and Law*, *7*, 97-113.

University of Utrecht (2005). *Intelligent Systems Group, Institute for Information & Computing Sciences Homepage*. Retrieved May 1, 2005, from http://www.cs.uu.nl/groups/IS/cd/cd.html

Verheij, B. (1996). *FAPR '96: The international conference on formal and applied practical reasoning* [unpublished conference report]. Retrieved May 1, 2005, from http://www.ai.rug.nl/~verheij/publications/fapr96report/

Walton, D., & Krabbe, E. C. W. (1995). *Commitment in dialogue: Basic concepts of interpersonal reasoning*. Albany: State University of New York Press.

Walton, D., & Godden, D. M. (forthcoming). Persuasion dialogue in online dispute resolution. *Artificial Intelligence and Law.*