Research Challenges for Argumentation
(Dagstuhl Perspectives Workshop 08042)[*]

Jürgen Dix
TU Clausthal, Germany[†]
Simon Parsons
Brooklyn College, United States of America
Henry Prakken
Utrecht University and University of Groningen, The Netherlands
Guillermo Simari
Universidad Nacional del Sur, Argentina

1st August 2008

Contents

---

1 Introduction

The first articles on argumentation in computer science appeared circa 20 years ago. Since then we have seen great advances, establishing a solid theoretical basis, a broad canvas of applications, and, most recently, some realistic implementations. The field has gone from infancy to maturity, and the initial questions that researchers posed—"how do we do this?", "what is it good for?" and "how do we implement it'?"—are mostly answered.

1.1 Argumentation as a Subarea in Computer Science

The problem of reasoning is at the core of Computer Science and was one of the early challenges that interested researchers such as Alan Turing, John McCarthy and Claude Shannon, among many others.

The first attempts to apply the entire apparatus of classical logic to solve the problem exposed some of its inadequacies and led to the exploration of different possibilities. Monotonicity is one of the properties of classical logic that was considered as a limiting feature when modeling commonsense reasoning. The formalization of systems that do not satisfy this property led to the introduction of *Nonmonotonic Reasoning Systems* in Artificial Intelligence (AI)

The more general concept of *Defeasible Reasoning* was introduced to the Computer Science part of AI in the mid 80's by the philosopher John Pollock [9],opening the field to more general alternatives. Research in Defeasible Reasoning in the AI area of Knowledge Representation and Reasoning has lead to important and useful results for other areas such as the development of intelligent agents and multi-agent systems (MAS) applications.

Modeling commonsense reasoning in AI is a difficult task given that it almost always occurs in the face of *incomplete* and potentially *inconsistent* information. Argumentation formalisms are defeasible reasoning systems which work by considering the reasons that lead to a given conclusion (or claim) through a piece of reasoning (the supporting argument) and the potential challenges (or counter-arguments) for accepting that conclusion. In this manner, the mechanisms proposed model reasoning as a *dialectical process*, *i.e.*, the exchange of arguments and counter-arguments respectively advocating and challenging the claim of the initial argument. This process offers a remarkable tolerance to the problems introduced by the potential inconsistency and/or incompleteness of the knowledge source.

Within the study of argumentation, we can distinguish a number of strands, each of which is an active area of research, and between which there is unfortunately little cross-fertilization. Argumentation's roots in models of disputation, for example, mean that it is well suited for the study of models of *dialogue*. Its ability to capture and resolve different points of view means that argumentation has found many applications in the area of *multi-agent systems*, since the different agents in such systems frequently have mutually inconsistent information and goals.

Researchers in *computational logic* have found ideas from argumentation helpful in establishing formal models of how idealized reasoning should proceed, and, similarly, techniques from argumentation have been used to *reason under uncertainty* and to implement models of *decision making.*

Finally, the ability of systems of argumentation to express and resolve complex forms of reasoning, especially reasoning which makes use of the interplay between conflicting, context-sensitive, rules, makes it a natural fit with the areas of *law* and *e-governance.*

1.2 Structure of the Perspectives Workshop

The group (28 participants from 12 countries) convened in Dagstuhl in January 2008, for a three day meeting.

The main objective of this Perspectives Workshop was to bring together researchers and developers involved in *Argumentation Systems and Related Areas* who have been producing significant contributions towards the progress in theoretical and pragmatic aspects of this form of reasoning.

During the preparation of the meeting the organizers decided to separate the group into four areas:

- **Argumentation and the Semantic Web**
  Speakers: Francesca Toni, Grigoris Antoniou
  Participants: Grigoris Antoniou, Philippe Besnard, Floriana Grasso, Gabriele Kern-Isberner, Thomas Meyer, Francesca Toni, Iyad Rahwan

- **Argumentation and Decision Support in Application Areas**
  Speakers: Anthony Hunter, Tom Gordon
  Participants: Pietro Baroni, Kevin Ashley, Trevor Bench-Capon, Tom Gordon, Anthony Hunter, Helmut Horacek, Guillermo Simari, Bart Verheij

- **Argumentation and Multi-Agent Systems**
  Speakers: Pavlos Moraitis, Leila Amgoud
  Participants: Leila Amgoud, Jürgen Dix, Paul E. Dunne, Antonis Kakas, Nicolas Maudet, Pavlos Moraitis

- **Argumentation and Social Networks**
  Speakers: Simon Buckingham Shum, Peter McBurney
  Participants: Simon Buckingham Shum, Peter McBurney, Tim Norman, Henry Prakken, Chris Reed, Simon D. Parsons, Douglas Walton

For each of these areas, two participants were requested to give talks to the plenary as a foundation for the discussions that followed.

After completing the presentations, the participants attended separate meetings according to the areas mentioned. The groups reported the conclusions to the plenary after each meeting.

During the last plenary, because of the results obtained so far and the enthusiasm of the participants, it was considered to produce publications for each of the four areas.

2 Looking Ahead

In this section we discuss for each of the above mentioned four areas (1) their relation to argumentation by analyzing the current situation and (2) we point out specific research challenges that we consider important and promising for future research.

2.1 Argumentation and the Semantic Web

An important aspect of creating ontologies, *Ontology Engineering*, is to allow for the collaborative development of ontologies, especially tracking the effect of design decisions,

and the ramifications of altering aspects of the ontology. The built-in dependency-tracking of argumentation systems makes them a natural technology to adopt.

The stumbling block in much work on ontology mapping and integration is that different ontologies define the same terms in ways that are directly or indirectly contradictory. Argumentation, being able to handle contradictions in a well-founded manner, can help to solve this problem of *Ontology interoperability* by supporting the identification and resolution of this contradictory information.

Finally, we would like to mention *personalisation*. As semantic web applications expand, a natural use of the semantic web will be *recommender systems*, especially systems that can reason from some specification provided by the user in order to identify suitable goods and services. Argumentation has proved a useful technology to support such applications as Chesñevar's ArgueNet system [4] shows.

2.1.1   The Semantic Web

The problem with the world wide web as it stands is that while there is an awful lot of content, that content is not very manageable. While this anarchy appeals to some, for anyone who uses the web as a tool—and that is an increasing number of us—this is a problem. Despite the power of services such as Google, any use of the web tends to involve a great deal of *manual* search that still requires additional effort from the user: Only very few *automated* services are available. The idea behind the semantic web is to improve the usability of the web by providing more automated services and thereby helping the user to extract the information that she wants.

The way that the semantic web aims to improve this usability is to allow web content to be augmented with a variety of meta data that explains what the content is and what the content does. A very early indication of the power of this kind of extension of basic content is provided by Technorati[1] tags, a mechanism for manually labelling content on blogs so that the Technorati blog tracking service can identify and link to relevant posts, providing a mechanism for easy identification of relevant material which is less error-prone than a keyword search through the blog entries.

The problem with Technorati tags, however, is that everyone who tags blog posts chooses their own tags — there is no common set of tags defined somewhere that one picks from — so when I tag an entry "CAT" (for the TAC market design competition) somebody can easily mistake this for an entry about small furry mammals. The semantic web deals with this problem through the use of ontologies, formal representations of the relationships between terms that enable them to be used in a consistent manner. If meta data is expressed using a vocabulary that is defined in an ontology, then the meta data has a defined meaning (which is where the "semantics" come in) and if the meta data is written in a formal language, then one can perform inference on it, and be able to interpret the results of inference.

The three pillars of the semantic web are thus:

- The use of meta data to annotate data on the web,

- the linking of the annotations to ontologies that relate the terms used in the meta data; and

- the use of logic-based techniques to reason over meta data and ontologies.

---

[1]http://www.technorati.com

Over the last few years, much progress has been made in defining the infrastructure that will support the semantic web, and a lot of effort has gone into defining a number of ontologies. However, there are still many open issues. For example, since a number of different ontologies have been developed, and there is no real way to say what is the best ontology (nor is it likely that all developers of semantic web applications would adopt the same ontology) it becomes important to be able to use several ontologies together. Hence *ontology mapping*, identifying how terms in one ontology relate to terms in another ontology, is an important issue. In addition, since the use of terminology changes over time, it becomes necessary to handle *ontology evolution* and *ontology change*. It turns out that argumentation, providing, as it does, flexible inference and the ability to recall why an inference was made, is helpful in a number of these tasks.

2.1.2 Research Challenges

In addition to the areas listed above where argumentation has been applied to solve problems in the provision of the semantic web, there are a number of places where argumentation might plausibly be used to handle issues with the semantic web that have not yet been considered in any detail.

Content Integration: One of the challenges of the semantic web that have yet to be addressed is the problem of combining information from different web repositories. Part of this problem is the problem of ontology interoperability that is discussed above. But even when ontologies are aligned, there are difficulties in using information drawn from different sources — there may be a need to resolve inconsistencies in the information, to draw plausible inferences to complete missing information, and to track the sources of data. Argumentation supports all these activities.

Information Aquisition: Right now constructing ontologies is a laborious, and hence time-consuming, business. Given the use of argumentation to support dialogue in multi-agent systems, it is easy to imagine that similar technology could be used to allow machine-human dialogues in which ontologies are automatically constructed, with arguments used to define terms, and then being used to automatically slot terms into the relevant places in an ontology.

Interactive Question Answering: Just as argumentation-based dialogue can be used to provide information that is used to construct ontologies, it can be used in semantic web applications to support user interaction.

2.2 Argumentation and Decision Support in Application Areas

A decision is a determination arrived at after consideration of all the relevant information. Decision making is therefore a defining component of any conceivable autonomous agent architecture, affecting not only the internal workings of the agent but its social behavior as well. Its paramount importance lies in the fact that once a decision is made, a course of action is taken that conditions the future of the agent and leaves behind other possible choices. Thus, decision making should be supported by reasoning that will account for the characteristics of the different alternatives that are available.

The cognitive process that supports the alternative, by its very nature, should be clear on the information used and the *way* it is used, *i.e.*, what is the informational base of the decision and how this information is processed in order to reach this decision. Furthermore, this process should involve an explicit consideration of alternate decisions,

and the reasons that support them being discarded should also be available to the agent. The characterization given applies to different entities that range from agents that decide what to do next as an individual to complex multi-agent systems that reach an agreement for their future behavior [7].

Argumentation-based reasoning is specially apt for reasoned decision making, since it satisfies the above mentioned conditions of providing all the elements necessary to understand the reasons behind a decision: The reasoning process supporting a decision and the possible considerations against it. This explanatory power, which distinguishes deliberative decision making in individuals, is particularly important when the decision should be made in the context of a multi-agent system where all the parties must accept the outcome of the process: Argumentation and its role in this area will be discussed in the next section.

Regarding decision making, the capability of presenting the reasons for and against a decision is of particular importance when the interaction among agents is considered. Every agent that is involved in making the decision can understand and challenge these reasons. This scenario will be addressed in the next section (Argumentation and MAS).

Software agents can present justifying reasons for a decision that the human counterpart can understand, and the human agent can revise her own reasoning process by considering lines of argumentation that were previously not taken into account. In this section, we only consider the single agent case.

2.2.1 Autonomous Agents

An agent becomes autonomous when, based on the state of the perceived environment, it possesses the capability of deciding what action to take. When the agent's decisions are based on an explicit consideration of the environment together with the internal state of the agent, its past experience, and its goals, the agent is considered *rational*. A significant part of human behavior can be modeled in this way, as well as the cognitive agents conceived in Artificial Intelligence research [10]. Much work has been done in trying to perfect models of reasoning based on argumentation [11, 3, 5]. A particularly strong point of an argumentation-based reasoning engine is that reasons for and against a given decision can be explicitly shown, opening them to further deliberation. This characteristic is especially useful when interacting with other entities [6].

As any reasoning system based on some form of logical language and some inference mechanism, argumentation systems face the issue of computational intractability. Some attempts have been made at improving the efficiency of the algorithms for building arguments and deciding on the warranting status of a given piece of knowledge. Most argumentation systems build dialectical structures to decide if a claim could be supported. There are several possible protocols that can be used to build the dialectical structure leading to different computational behaviors (skeptic vs. credulous, finite structures, etc.) Certainly, these are lines of research that are worthwhile to pursue.

2.2.2 Research Challenges

More research is needed in the areas of *argument construction, argument evaluation, argument visualization, dialogue and protocol managers*, and *argument presentation and extraction tools.*

In particular, we consider the following branches especially promising and worthwhile to consider.

Extending: What does an argumentation-theoretic approach add over and above decision theory? How can one integrate argumentation tools with classical decision theory and other existing models of decision making?

Symbolic vs. non-symbolic: What role should argumentation systems play in integrating symbolic and non-symbolic decision making methods? How should one connect non-symbolic systems such as neural networks to argumentation tools?

Relations to machine learning: What are the connections between case-based argumentation schemes, decision support, and machine learning? How can one incorporate data stores into the argumentation tools using machine learning and data mining?

Integration with databases: How can one use arguments to make decisions with large amounts of data?

What role can argumentation systems play in teaching novices how to make decisions?

2.3 Argumentation and Multi-Agent Systems

Multi-agent systems provide a fertile area for the application of argumentation-based techniques. There are specialized workshops (e.g ArgMAS within AAMAS, the most influential conference on MAS) and even stand-alone conferences (COMMA) that receive contributions in many different areas.

From *Agent Communication* to *Agent Societies*, there is a wide range of open problems that can be considered from the point of view of argumentation. Several of the issues corresponding to multi-agent systems are addressed in this section.

In the multi-agent systems area, the presentation of arguments in such a way that can be understood by humans and other types of agents is an important line of research that involves a movement between the artificial agent formal side and the human, less formal, presentation. Related to this, the development of argumentation tools for mixing human and software reasoning will become increasingly necessary. Dynamic scenarios pose a particularly tough problem. Here, research on *anytime algorithms* for deciding on the status of warrant is necessary.

Argumentation can be seen as the interaction of different arguments for and against some conclusion. In the last 10 years, argumentation has been gaining increasing importance in multi-agent systems, mainly as a vehicle for interaction which involves the *giving and receiving of reasons*. Argumentation has made solid contributions to the practice of *multi-agent dialogues*, such as *legal disputes*, *business negotiation*, *coalition formation*, *ontology reconciliation*, *risk analysis*, *scheduling*, and *logistics*. A single agent may also use argumentation techniques to perform its individual reasoning because it needs to make decisions under complex preferences policies, in a highly dynamic environment.

2.3.1 The Multi-Agent Case

An important task in multi-agent systems is *negotiation*. This is a form of interaction in which autonomous agents (with different interests and goals) try to find a compromise on a certain issue. Further well-known tasks in MAS are (1) *collaboration* (delegating a goal or task to another agent, (2) *communicating* with other agents, (3) *coalition formation* (joining a group of agents if that is beneficial for the agent and solving the overall task).

This leads to the study of the *properties of protocols* allowing to reach optimal solutions: How can an optimal solution be reached quickly? Can a meta model for decision making be defined to chose the next move to play in a dialogue?

The most important component within an agent is the *reasoning component*. It allows the agent to *deliberate* about its current situation and decide which actions to take. Often, such reasoning requires large amounts of heterogenous data. Therefore, we need reasoning mechanisms that are suitably *resource-bounded*: Justified conclusions should be computed in reasonable time.

A major issue in MAS applications is the issue of *trust*: Who and what is it reasonable to rely on? Argumentation has been used in a number of applications to handle trust, for example weighting inferences with measures of trust in the source of the information, and updating these measures as sources prove to be more or less reliable. This matter has been considered at some length by the argumentation community, and the resulting ideas, such as Pollock's notion of warrant [9],can be used to provide a solution to the issue.

2.3.2 Research Challenges

Core Engine: ARG has been used for many important tasks in MAS (negotiation, coalition-formation, persuasion, deliberation). Can argumentation be the *core engine of an agent*, instead of the mental state? This would allow a seamless integration of these tasks and provide a unified view of the reasoning component of an agent.

Trust: Which agents are *trustworthy*? This is important for taking decisions and weighing arguments of other agents. ARG has already been applied to *weigh inferences* wrt. trust in sources given as databases. Can this theory be extended to *define a notion of trust between agents*?

Tractability: In almost all theories of argumentation, the problem of deciding on the *status of a warrant* is paramount. Therefore this problem is at the heart of the overall complexity of argumentation systems. Are there *anytime algorithms* for this task? This would help for

2.4 Argumentation and Social Networks

Virtual social networks exist in many forms: small or large, open or closed, consisting of the general public, specialists or a mixture; collaborative or adversarial; outcome-oriented or not; with emphasis on roles and sources of opinion or not. The subgroup on social networks has identified a number of use cases for argumentation in virtual social networks, illustrating different points at these dimensions. For each use case the current state of argumentation support has been determined and the main research challenges identified. Below we discuss and contrast two of these scenarios, namely, online political debate and crime analysis. These scenarios are different in many respects and are therefore very suited for illustrating the main issues and challenges. The other scenarios are: e-learning, scientific publishing, law firm negotiation, business forecasting, and argument for fun.

However, for now one common feature of all these use case scenarios must be mentioned. Generally in social networks the arguments will be "pulled from" the user rather than "pushed onto" them by a knowledge-based automated reasoner (cf. Hunter [8]). This means that argumentation applications in social networks will look less at how machines can produce argumentation from represented information and more at how humans can

be supported in structuring their arguments and debates so that machines can utilise such structure to provide further services to the humans.

2.4.1 Two Use Case Scenarios

Online political debate Online communities that are engaged in political debate typically are open and consist of a mixture of the general public and specialists (for example journalists, lobby groups, policy analysts). The setting typically is adversarial. For most participants the outcome is important but there is no overall necessity to reach a consensus or decision. Roles generally are not very important (politicians hardly engage in these debates, presumably since they see no advantage in doing so) but people often back their opinions with evidence or support by e.g. linking to documents, other websites, pictures or videos.

Currently, online political debate largely occurs in general media, such as discussion fora and blogs. In most cases the only structuring of the debate is that of threaded discussions but a few websites support more structuring. A failed experiment was the Dutch site www.democratie-nu.org, at which the Dutch general public could discuss political issues in a format of nested arguments pro and con; the idea was that politicians would regularly check the site to obtain information on the citizens' opinions. This site was launched with much publicity but it was never really used, probably because of a poor interface and lack of tools for summarising and overview. A current site with (non-nested) pro and con structure is `www.debatepedia.com` (not limited to political debate). However, there still is a considerable gap between the structured tools and theories developed in research and what practice wants and is able to.

Given that the general public wants easy-to-use tools the question arises why we should aim for more structure than in current practice. Still there are good reasons to do so. For instance, structured arguments and debates can be used by tools for generating overviews and syntheses of a debate, and they are more likely to be picked up by aggregation services. Therefore, future research should focus on easy-to-use tools for structured individual argument authoring and publishing. Since Web 2.0 users more and more refer to movies, pictures and so on, media annotation tools are also an important research topic.

A major research challenge here is how users can be stimulated or seduced to structure their contributions. One thing to profit from is that in large open social networks an important driving factor is *reputation* and the most respected discussion members are those who can summarise and reconstruct the debate and list the options. Therefore, future research should also focus on tools for annotating, summarising, and reconstructing a debate. Since the "sense makers" using these tools may be specialists or at least willing to spend more time on learning tools, the tools can be more specialised and advanced than those for individual argument authoring. An even further step is to (semi-)automate such tools; this requires an integration with natural-language processing and advanced search techniques.

To summarise, what is needed is easy-to-use light-weight argument authoring and annotation tools for the general public and more advanced debate structuring tools for self-appointed 'sense makers'. Ideally, such tools can be integrated with existing social-network software on the web, such as discussion fora, blogs, and websites such as YouTube or Facebook.

Crime analysis The crime analysis scenario (and similar scenarios such as intelligence analysis) concerns a small team of crime analysts in police forces, whose task is to analyse

a mass of evidence related to one or more purported crimes, to construct and investigate hypotheses of what might have happened and to choose one or more lines of further investigation. This scenario is an example of a small and closed social network with collaborating professionals as members, who employ specific and advanced reasoning styles. Roles are not so important but the sources of arguments are extremely important. These sources often come in the form of physical artifacts, multi-media or links with off- or online data sources of various kinds (telephone directories, fingerprint or DNA databases, Google maps, and so on). The outcome of the crime investigation process is, of course, all important.

The need of crime analysts for software tools for structuring their thinking and discussions is well-documented. However, current specialised crime analysis software (such as Analyst's Notebook) typically only allows them to structure the available evidence into ontology- and time-line like structures, by expressing relations of various kinds between persons, objects and events. The process of formulating hypotheses and assessing their relation with the evidence is left to the human analyst and the structures resulting from such reasoning cannot be recorded and analysed by the software, so that it cannot be communicated to other team members. Recent research aims at overcoming these limitations. For example, Bex et al. [2] report on a structuring tool for combined argumentation and abductive reasoning. However, much more needs to be investigated. Prime research issues are the annotation of evidence in multi-media form and the linking of argument maps with the available databases of various kinds.

Since the crime analysis task is a specific one carried out by professionals, specific forms of reasoning are used so the support should also be specialised and advanced. Acceptance of tools imposing structure is more likely than with the general public, since professionals might be willing to invest more learning time, or this might be part of their education and training. However, to make the tools effective, they must be based on a thorough analysis of the work practices of the professionals involved.

2.4.2 Research Challenges

In social networks the arguments will be "pulled from" the user rather than "pushed onto" them by a machine. Therefore it seems little place in social networks research for 'traditional' AI-style knowledge representation and reasoning. At best, automated reasoners seem useful for evaluating the outcome of debates given the arguments and preferences provided by humans; to yield interesting results, this requires highly structured arguments and debates, so automated evaluation tools seem mainly relevant for small-scale and professional social networks where much depends on the outcome, such as the crime investigation scenario.

Automated Tools: How can users be "convinced" to structure their contributions in a way that is manageable for them while it enables the use of *advanced automated argumentation tools*?

Argument Authoring: The importance of tools for multi-media annotation in argument authoring has been shown. *What is the structure and role of visual argumentation?* Advanced, easy to use debate structuring tools are needed.

Use Cases: *Online political debate* and *Crime analysis*. Structure not only available evidence (*Analysts Notebook*) but also the *process of formulating hypotheses*, linking it with other sources etc.

3   Conclusion

Argumentation systems have reached a stage of theoretical maturity where the possibilities for technological innovation and industrial applications have become a reality (see [6, 3, 1]. Leading research groups have produced prototypical software systems with argumentation-based reasoning engines as proof-of-concept. Such implementations are not only research tools but they have become concrete options.

The EU-funded ASPIC project[2] (on which H. Prakken and S. Parsons have been working) has developed industrial-strength Java components that implement an argumentation system, and which will be made available under an open-source license. These components will make it possible to construct software systems that make use of the power of argumentation in the service of other functionality (rather than as an end in itself).

Another EU-funded program which aims at using the aforementioned components is the *ArguGRID* project[3], on which T. Kakas and F. Toni are working on. Argumentation technology is used to support rational decision making. Agents use this as an internal mechanism and for supporting negotiation with other agents in the dynamic composition of Grid services/resources into an executable workflow. Argumentation is also used to support the creation, management and dynamic evolution of societies of agents to compose complex services from individual ones.

These are just two well-known examples of technological innovation. Several other systems based on different approaches have been developed by other groups following other promising avenues, *e.g.* the *Semantic Web and Cognitive Robotics*.

At the same time, the research community associated with argumentation has been growing steadily during the last decade. The existence of a full conference dedicated to the topic, along with several specialized workshops on different areas of argumentation, shows a growing interest in the AI community which is very promising. In particular, in May 2008, during the Second International Conference on Computational Models of Argument, there was a session dedicated to Software Demonstrations which showcased some of the most recent developments.

The field is now ready for the next iteration of the research cycle in which the proof-of-concept implementations are refined and their usage produce new problems for basic research whose solution will enhance the systems' robustness and functionality. We believe that the application of argumentation tools is gaining significance in a growing number of different areas, *e.g.* Human-Computer Interaction and Social Networks: Not just as theoretical tools, but as practical systems that can be incorporated into other systems.

A   Participants

Leila Amgoud	IRIT - Toulouse, FR
Grigoris Antoniou	FORTH - Heraklion, GR
Kevin Ashley	University of Pittsburgh, US
Pietro Baroni	University of Brescia, IT
Trevor Bench-Capon	University of Liverpool, UK
Philippe Besnard	Université de Toulouse, FR
Simon Buckingham Shum	The Open University - Milton Keynes, UK

---

[2]http://www.argumentation.org/
[3]http://www.argugrid.eu/

Jürgen Dix	TU Clausthal, DE
Paul Dunne	University of Liverpool, UK
Tom Gordon	Fraunhofer FOKUS, DE
Floriana Grasso	University of Liverpool, UK
Helmut Horacek	Saarland University, DE
Anthony Hunter	University College London, UK
Antonis Kakas	University of Cyprus, CY
Gabriele Kern-Isberner	Universität Dortmund, DE
Nicolas Maudet	Université Paris-Dauphine, FR
Peter McBurney	University of Liverpool, UK
Thomas Meyer	Meraka Institute - Pretoria, ZA
Pavlos Moraitis	Paris Descartes Université, FR
Tim Norman	University of Aberdeen, UK
Simon Parsons	Brooklyn College, US
Henry Prakken	Utrecht University, NE
Iyad Rahwan	British University in Dubai, UAE
Chris Reed	University of Dundee, UK
Guillermo Simari	Universidad Nacional del Sur, AR
Francesca Toni	Imperial College London, UK
Bart Verheij	University of Groningen, NE
Doug Walton	University of Winnipeg, CA

B   References

[1]   Ph. Besnard and A. Hunter. Elements of argumentation. MIT Press, 2008.

[2]   F.J. Bex, S.W. van den Braak, H. van Oostendorp, H. Prakken, B. Verheij, and G. Vreeswijk. Sense-making software for crime investigation: how to combine stories and arguments? *Law, Probability and Risk*, 6:145–168, 2007.

[3]   C. Chesñevar, A. Maguitman and R. Loui. Logical Models of Argument. ACM Computing Surveys. Vol 32, no 4, pp 337–383, 2000.

[4]   C. I. Chesñevar and A. G. Maguitman. ArgueNet: An Argument-Based Recommender System for Solving Web Search Queries in *Second IEEE Conference on Intelligent Systems*, 2004.

[5]   Trevor J. M. Bench-Capon, Paul E. Dunne. Argumentation in artificial intelligence. Artificial Intelligence, 171(10-15), pp 619–641, 2007.

[6]   Alejandro J. Garcia, Nicolas D. Rotstein, Guillermo R. Simari. Dialectical Explanations in Defeasible Argumentation. Procs. of the 9th European Conf on Symbolic and Qualitative Aspects of Reasoning with Uncertainty (ECSQARU 2007). LNCS v.4724, Pp. 295—307. Tunez. 2007.

[7]   R. Girle, D. Hitchcock, P. McBurney, and B. Verheij. Decision Support for Practical Reasoning in *Argumentation Machines: New Frontiers in Argument and Computation*, C. Reed and T. Norman (Eds.), Chapter 3, pp 55–84. Kluwer Academic Pub. 2004.

[8]   A. Hunter. Hybrid argumentation systems for structured news reports. *The Knowledge Engineering Review*, 16:295–223, 2001.

[9]   J. Pollock. Defeasible Reasoning. *Cognitive Science*, 11, pp 481–518, 1987.

[10]   J. Pollock. Against optimality: Logical Foundations for Decision-Theoretic Planning in Autonomous Agents. *Computational Intelligence*, 22:1, pp 1–25, 2006.

[11]   H. Prakken and G. Vreeswijk. Logical Systems for Defeasible Argumentation. Handbook of Philosophical Logic, D. Gabbay and F.Guenther (Eds.), Kluwer Academic Publishers, pp 219–318, 2002.