Argumentation in Multi-Agent Systems: Context and Recent Developments

Nicolas Maudet¹, Simon Parsons², and Iyad Rahwan³

¹ LAMSADE, Université Paris-Dauphine
Paris 75775 Cedex 16, France
maudet@lamsade.dauphine.fr

² Department of Computer and Information Science, Brooklyn College
City University of New York, 2900 Bedford Avenue, Brooklyn, 11210 NY, USA
parsons@sci.brooklyn.cuny.edu

³ Institute of Informatics, The British University in Dubai
P.O.Box 502216, Dubai, UAE
(Fellow) School of Informatics, University of Edinburgh, UK
irahwan@acm.org

Abstract. This chapter provides a brief survey of argumentation in multi-agent systems. It is not only brief, but rather idiosyncratic, and focuses on the areas of research that most interest the authors, and those which seem to be the most active at the time of writing.

1 Introduction

The theory of argumentation [81] is a rich, interdisciplinary area of research lying across philosophy, communication studies, linguistics, and psychology. Its techniques and results have found a wide range of applications in both theoretical and practical branches of artificial intelligence and computer science [14,74]. These applications range from specifying semantics for logic programs [20], to natural language text generation [21], to supporting legal reasoning [9], to decision-support for multi-party human decision-making [31] and conflict resolution [80].

In recent years, argumentation theory has been gaining increasing interest in the multi-agent systems (MAS) research community. On one hand, argumentation-based techniques can be used to specify *autonomous agent reasoning*, such as belief revision and decision-making under uncertainty and non-standard preference policies. On the other hand, argumentation can also be used as a vehicle for facilitating *multi-agent interaction*, because argumentation naturally provides tools for designing, implementing and analysing sophisticated forms of interaction among rational agents. Argumentation has made solid contributions to the theory and practice of multi-agent dialogues.

In this short survey, we review the most significant and recent advances in the field, with no intention of being exhaustive. Thus, we ignore recent work that extends the basic mechanisms of argumentation with new semantics [12], bipolar arguments [13], and the ability to handle sets of arguments [49]. Indeed, we have

very little to say about *how to argue* and, instead, deal with *what one can argue about*, dealing with the uses of argumentation rather than the mechanisms by which it may be carried out¹, and restricting even that view to coincide with the topics of the other papers in this volume. In particular, this chapter first recalls some of the key notions in argumentation theory, and then outlines work on two major applications of argumentation in multi-agent systems, namely in the reasoning carried out by autonomous agents (Section 3) and in multi-agent communication (Section 4).

2 What Is Argumentation Good for?

According to a recent authoritative reference on argumentation theory, argumentation can be defined as follows:

Argumentation is a verbal and social activity of reason aimed at increasing (or decreasing) the acceptability of a controversial standpoint for the listener or reader, by putting forward a constellation of propositions intended to justify (or refute) the standpoint before a rational judge. [81, page 5]

Let us decompose the elements of this definition that are most relevant to our discussion. First, the ultimate goal of argumentation is to resolve a “*controversial*” standpoint; controversial in the sense that it is subject to both “*justification*” or “*refutation*” depending on the information available. This distinguishes argumentation from the classical deductive reasoning viewpoint, in which proofs for propositions cannot be contested. Moreover, the nature of the “*standpoint*” can vary. While the classical study of argumentation has focused mainly on propositional standpoints — i.e. things that are believed or known — there is no reason why the standpoint is confined to be propositional. A standpoint can, in principle, range from a proposition to believe, to a goal to try to achieve, to a value to try to promote. That is, argumentation can be used for theoretical reasoning (about what to believe) as well as practical reasoning (about what to do).

Secondly, argumentation is an “*activity of reason*”, emphasising that a particular process is to be followed in order to influence the acceptability of the controversial standpoint. This activity and the propositions put forward are to be evaluated by a “*rational judge*”: a system that defines the reasonableness of these propositions according to some criteria. An important objective of argumentation theory is to identify such system of criteria.

In summary, argumentation can be seen as the principled interaction of different, potentially conflicting arguments, for the sake of arriving at a consistent conclusion. Perhaps the most crucial aspect of argumentation is the interaction between arguments. Argumentation can give us means for allowing an agent to

---

¹ Not least because one can potentially make use of any mechanism for argumentation in the service of any of the applications of argumentation.

reconcile conflicting information within itself, for reconciling its informational state with new perceptions from the environment, and for reconciling conflicting information between multiple agents through communication. It is for these reasons that argumentation has begun to receive great interest in the multi-agent systems community. In particular, argumentation lends itself naturally to two main sorts of problems encountered in MAS:

- **Forming and revising beliefs and decisions:** Argumentation provides means for forming beliefs and decisions on the basis of incomplete, conflicting or uncertain information. This is because argumentation provides a systematic means for resolving conflicts among different arguments and arriving at consistent, well-supported standpoints;
- **Rational interaction:** Argumentation provides means for structuring dialogue between participants that have potentially conflicting viewpoints. In particular, argumentation provides a framework for ensuring that interaction respects certain principles (e.g. consistency of each participant’s statements).

In the next sections, we will discuss these applications in more detail and refer to some relevant literature. In particular, Section 3 deals with the topics of revising beliefs and making decisions, aspects that we can think of as being the concern of individual autonomous agents, while Section 4 deals with topics related to inter-agent communication and rational action, all aspects of argumentation that are decidedly multi-agent.

3 Argumentation for Reasoning in Autonomous Agents

Argumentation is a general process for reasoning. An autonomous agent that has to reason about could weigh arguments for and against different options in order to arrive at a well-supported stance. In this section, we discuss two main applications of argumentation to autonomous agent reasoning.

3.1 Argumentation for Belief Revision

One of the main challenges in specifying autonomous agents is the maintenance and updating of its beliefs in a dynamic environment. An agent may receive perceptual information that is inconsistent with its view of the world, in which case the agent needs to update its beliefs in order to maintain consistency. The major challenge of nonmonotonic reasoning formalisms [11] is to specify efficient ways to update beliefs. At the normative level, the AGM paradigm [29] specifies the rationality postulates that must be satisfied by an idealistic process of belief revision. On the operational level, formalisms for mechanising nonmonotonic reasoning include truth maintenance systems (TMS) [19], default logic [75] and circumscription [48].

Argumentation provides an alternative way to mechanise nonmonotonic reasoning. Argument-based frameworks view the problem of nonmonotonic reasoning as a process in which arguments for and against certain conclusions

are constructed and compared. Nonmonotonicity arises from the fact that new premises may enable the construction of new arguments to support new beliefs, or stronger counterarguments against existing beliefs. As the number of premises grows, the set of arguments that can be constructed from those premises grows monotonically. However, because new arguments may overturn existing beliefs, the set of beliefs is nonmonotonic. Various argument-based frameworks for non-monotonic reasoning have been proposed in the last 20 or so years. Some of the most notable are the following [42,60,79,41,22,27,67]².

While the above-mentioned frameworks have developed into a solid and mature sub-field of AI, their incorporation into *situated* autonomous agent reasoning remains an opportunity to be pursued. In order to do so, an adequate representation of the environment is needed, and a mechanism for integrating perceptual information into the belief-update mechanism is also required. Moreover, situated agents are required to update their beliefs in a timely fashion in order to take appropriate action accordingly.

3.2 Argumentation for Deliberation and Means-Ends Reasoning

An autonomous agent does not only maintain a mental picture of its environment. The agent is faced with two additional tasks: the task of *deliberation* in which it decides what state of the world it wishes to achieve — namely its goal — and the task of *means-ends reasoning* in which it forms a plan to achieve this goal. Argumentation is also potentially useful for tackling both these challenges.

Recently, argumentation has been applied to deliberation. For example, argumentation has been used as a means for choosing among a set of conflicting desires [1] and as a means for choosing between goals [3]. Another argument-based framework for deliberation has been presented by Kakas and Moraitis [39]. In this approach, arguments and preferences among them are used in order to generate goals based on a changing context. In addition, argumentation can be used to support standard BDI [73] models, as in [56].

More generally, as shown by Fox in his work since [26]³, argumentation provides a framework for making decisions. Just as one makes arguments and counterarguments for beliefs, one can make arguments and counterarguments for actions. While such a framework sounds as though it must be at odds with approaches based on decision theory [34], Fox and Parsons [28] provide an argumentation framework that reconciles the two approaches. In this system, argumentation is used to reason about the expected value of possible actions. In particular, one argument system is used to arrive at a stance on beliefs, while another argument system identifies the *outcomes* of possible actions. Together, arguments over beliefs and the results of actions can be combined to create arguments about the *expected value* of possible actions. This approach was later refined in [53].

² For comprehensive surveys on argument-based approaches to nonmonotonic reasoning, see [14,68].

³ Though this line of work, summarised in [52], did not explicitly use the term “argumentation” until [27], with hindsight it is clear that argumentation is exactly what Fox and his colleagues were using.

Argumentation has also been used in planning. One of the earliest works on argument-based planning is perhaps George Ferguson's thesis [23], which uses argumentation as a means of allowing several participants to collaborate on the creation of a plan — plans are presented as arguments that a given course of action will result in a goal being achieved. Around the same time, John Pollock's was extending his OSCAR system to deal with the notion of defeat among plans [61]. More recently, several researchers have considered using argument-based approaches to generate plans [3,36,78]. However, such frameworks currently generate relatively simple plans in comparison with algorithms found in the mainstream planning literature [30]. One important question worth exploring is whether argumentation will offer real advances over existing planning algorithms.

4 Argumentation for Agent Communication

An inherent, almost defining, characteristic of multi-agent systems is that agents need to communicate in order to achieve their individual or collective aims. Argumentation theory has been an inspiration for studying and formalising various aspects of agent communication. Enhancing agent communication with argumentation allows agents to exchange arguments, to justify their stance, to provide reasons that defend their claims. This improved expressivity has many potential benefits, but it is often claimed that it should in particular:

- make communication *more efficient* by allowing agents to reveal relevant pieces of information when it is required during a conversation;
- allow for a *verifiable semantics* based on the agents' ability to justify their claims (and not on private mental states); and
- make protocols *more flexible*, by replacing traditional protocol-based regulation by more sophisticated mechanics based on commitments.

On the other hand, this improved expressivity comes with a price: it poses some serious challenges when it comes to designing autonomous agents that actually communicate by means of arguments, and makes more difficult:

- the *integration with agents' reasoning*, which requires to precisely specify what agents should respond to others' agents on the basis of their internal state, but also on the basis of their goal (strategy);
- the validation of *provable desirable properties* of these protocols;
- the communication between potentially *heterogeneous agents*, which should now share an *argument interagent format*.

We now critically discuss some of the points listed above, by questioning whether these hopes have been justified, and whether the aforementioned difficulties have seen some significant advances in recent years.

4.1 Efficiency of Argumentation

Until rather recently it was often claimed that argumentation could make communication more efficient, by allowing agents to reveal relevant pieces of information when it is needed during a conversation. Although the idea is intuitively appealing, there was little evidence to confirm this though. Indeed, argumentation may also involve both computational and communication overload, hence compensating the potential benefits induced by the exchange of reasons justifying agents' stances regarding an issue.

Perhaps the pioneering work in this area is that of Jung et al. [37,38]: in the context of a (real world) applications modeled as distributed constraint-satisfaction problems (e.g. a distributed sensor domain), they study whether the overhead of argumentation is justified by comparing various strategies. In the first edition of the ArgMAS workshop series, Karunatilake and Jennings [40] ask the question directly: "Is it worth arguing?". In the context of a task-allocation problem, they investigated how argumentative strategies compare to alternative means of resolving conflicts (evading, or re-planning). More recently, the efficiency of argument-based communication has been explored in the different context of a crisis situation involving agents trying to escape a burning building [10]. If agents make uncertain hypotheses regarding the origin of the fire, when should they waste time in trying to convince their partners? In this volume, Ontañón and Plaza [50] experimentally examine how argumentation can make multiagent learning more efficient.

Without entering in the details of these experimental results, it is interesting to note that the efficiency of argumentation is very much dependent of the context, and that there can be no straightforward answer to the question "Is it worth arguing?". For instance, Karunatilake and Jennings show that argumentation turns out to be effective when the the number of resources involved in the task allocation problem remains rather limited. Similarly, Bourgne et al. [10] observe that argumentation is especially required in those situations where buildings are rather "open" (when there are fewer walls). This can be explained by the fact that there are more potential candidate hypotheses to the fire origin then, hence the need to exchange arguments to discriminate between those. In their multi-agent learning experiment [50], Ontañón and Plaza emphasize in particular the influence of the amount of data that agents can individually access: as expected, argumentation is more beneficial when agents have only limited access to data.

While these papers try to investigate mostly experimentally in what circumstances argumentation can be an efficient conflict resolution technique; there are more theoretical contributions to this issue. A recent paper by Rahwan et al. [70] makes a first effort in this direction. In particular, the authors investigate a simple argumentation-based negotiation protocol in which agents exchange information about their underlying goals. It is shown that under certain conditions, exchanging such information enables agents to discover mutual goals and thus increases the likelihood of reaching deals. Other related work is that of [57] which shows how the beliefs of two agents that engage in argumentation-based dialogue will converge over time.

4.2 Flexibility of Communication

One of the most formally precise ways of studying different types of dialogues is through *dialogue-games*. Dialogue-games are interactions between two or more players, where each player makes a *move* by making some utterance in a common communication language, and according to some pre-defined rules. Dialogue-games have their roots in the philosophy of argumentation [7] and were used as a tool for analysing fallacious arguments [32]. Such games have been used by Walton and Krabbe themselves to study fallacies in persuasion dialogues.

Recently, dialogue-games have become influential in AI and MAS, mainly as a means for specifying protocols [44]. A *dialogue-game protocol* is defined in terms of a set of locutions, as well as different types of rules: *commencement rules*, *combination rules*, *commitment rules* and *termination rules* [46]. Commencement and termination rules specify when a dialogue commences and how it terminates. Commitment rules specify how the contents of commitment stores change as a result of different locutions. Finally, combination rules specify the legal sequences of dialogue moves.

In AI and MAS, formal dialogue-game protocols have been presented for different atomic dialogue types in the typology of Walton and Krabbe described above. These include persuasion dialogues [5], inquiry dialogues [35], negotiation [47,77], and deliberation [33]. Other types of dialogues based on combinations of such atomic dialogues have also been proposed, including team formation dialogues [17], dialogues for reaching collective intentions [18], and dialogues for interest-based negotiation [69].

Dialogue-game protocols offer a number of advantages. Mainly, they offer an intuitive approach to defining protocols and naturally lend themselves to argumentation-theoretic analysis, e.g. of dialogue embedding, commitments and fallacies. It is then feasible to define protocols that would otherwise be difficult to specify in practice, were we to use a different means of representation, for instance finite state machines (although their expressive power may not be higher in theory [24]). In practice, dialogue-games seem to offer a good compromise between the strict rule-governed nature of many implemented agent systems (economic auction mechanisms [84] being a good example) and the greater expressiveness envisioned by generic agent communication languages such as FIPA-ACL [25] (see [46]).

Now finding the good degree of flexibility is a difficult exercise. Designing the rules of a protocol amounts to specify what counts as a legal conversation between agents involved in a given interaction. Of course, the objective is to reduce the autonomy of agents in order to be able to prove interesting properties (see below), but at the time to allow agents to exchange arguments in a way that is deemed “natural” and flexible. For instance, the traditional proof-theoretical concept that takes the form of a dialectical dialogue between a proponent and an opponent can hardly be regarded as flexible: agents are highly constrained in their possible responses, with no possibly, for instance, to get back to a previous claim and explore alternative replies. In some circumstances (as was already

argued in [43]), it can be appropriate to leave agents explore the space of possible alternatives more widely. The work of Prakken has certainly been pioneering in this respect, in trying to articulate both the necessity of flexible protocols with concrete mechanisms while still maintaining their coherence [62,63,64]. The notion of *relevance* has been put forward as a central notion by Prakken: in very broad terms, moves are deemed legal when they respond to some previous move of the dialogue and are *relevant*, in the sense that they modify the current winning position of the dialogue.

4.3 Integration of Argumentation and Reasoning

We have seen that argumentation can serve both as a framework for implementing autonomous agent reasoning (e.g. about beliefs and actions) and as a means to structure communication among agents. As a result, argumentation can naturally provide a means for integrating communication with reasoning in a unified framework.

To illustrate the above point, consider the following popular example by Parsons et al. [56]. The example concerns two home-improvement agents — agent A_1 trying to hang a painting, and another A_2 trying to hang a mirror. A_1 possesses a screw, a screw driver and a hammer, but needs a nail in addition to the hammer to hang the painting. On the other hand, A_2 possesses a nail, and believes that to hang the mirror, it needs a hammer in addition to the nail. Now, consider the following dialogue (described here in natural language) between the two agents:

A1: *Can you please give me a nail?*

A2: *Sorry, I need it for hanging a mirror.*

A1: *But you can use a screw and a screw driver to hang the mirror! And if you ask me, I can provide you with these.*

A2: *Really? I guess in that case, I do not need the nail. Here you go.*

A1: *Thanks.*

At first, A_2 was not willing to give away the nail because it needed it to achieve its goal. But after finding out the reason for rejection, A_1 managed to persuade A_2 to give away the nail by providing an alternative plan for achieving the latter's goal.

We can use this example to highlight how argumentation-based techniques can provide a comprehensive set of features required for communication. Let us consider these in detail.

1. **Reasoning and Planning:** Argumentation can be used by each agent to form its beliefs about the environment, and to generate plans for achieving their goals. For example, agent A_2 can use argument-based deliberation to arrive at the goal to acquire a nail.
2. **Generating Utterances:** Argumentation can be used to generate arguments for utterances and arguments. For example, after A_1 requests a nail from A_2 , the latter builds an argument against giving away the nail by stating that it needs the nail to achieve one of its own goals (namely, hanging

the mirror). This information can be used again by A_2 to generate a counter-argument for why A_2 does not need the nail.

3. **Evaluating incoming communication:** Argumentation-based belief revision can be used to evaluate incoming communication. For example, when A_2 received the argument from A_1 , it had to evaluate that argument to make sure it is sensible. A_2 would not have accepted A_1 's argument if the former did not believe the latter actually possesses a screw and screw driver.
4. **Communication Structuring:** The whole dialogue can be structured through argumentation-based protocols, based on dialogue-games, which may themselves be based on certain argumentation schemes for reasoning about resources and plans.

Indeed, the above example, described in a theoretical framework by Parsons et al. [56], has been fully implemented using an argumentation framework based on abductive logic programming [77]. Other attempts to integrate reasoning and communication within a unified argumentation framework have also been made [6,76,69]. A review of these frameworks and others can be found in [71].

A major inspiration from argumentation theory in MAS is the notion of an *argumentation scheme* [83]. These are schemes that capture stereotypical (deductive or non-deductive) patterns of reasoning found in everyday discourse. For example, Walton specifies twenty five argumentation schemes for common types of presumptive reasoning. The most useful aspect of argumentation schemes is that they each have an associated set of *critical questions*. These critical questions help identify various arguments that can be presented in relation to a claim based on the given scheme. Hence, while a scheme can be used to establish a "stance," the set of critical questions help build communication structures about that stance.

Argumentation schemes offer a number of useful features to MAS communication. Their structure helps reduce the computational cost of argument generation, since only certain types of propositions need to be established. This very feature also reduces the cost of evaluating arguments.

A few attempts have been made to utilise the power of argumentation schemes in AI, mainly in constructing argumentation schemes for legal reasoning [82,66]. In MAS, the paper by Atkinson et al. [8] uses an argumentation scheme for proposing actions to structure their dialogue-game protocol.

A particularly important issue on the boundary between communication and internal reasoning is the specification of *argumentation dialogue strategies*. A strategy in an argumentation dialogue specifies what utterances to make in order to bring about some desired outcome (e.g. to persuade the counterpart to perform a particular action). While work on argument evaluation and generation has received much attention, the strategic use of arguments has received little attention in the literature. Recently, the effects of a specific set of agent *attitudes* on dialogue outcomes have been studied [4,59]. For example, a *confident* agent is happy to assert statements for which it has an argument, but a more *careful* agent makes assertions only after going through its whole knowledge base and making sure it has no arguments against it. When it comes to more complex

dialogue strategies, however, only informal methodologies have been proposed [69, Chapter 5].

Work on an agent’s strategy overlaps with the notion of relevance that was mentioned above. In a dialogue, unless it is very constrained, an agent typically has a choice of possible utterances. How the agent makes the choice is an aspect of its strategy, and relevance may come into its strategic thought. For example, as Oren et al. consider [51], an agent may be wise to avoid saying anything that is essential to the case it is making, for fear that it may be used against it at a later point⁴. Oren et al. use a notion of what is relevant, similar to that used by Prakken, to establish what an agent might sensibly say, and Bentahar et al. [45] make use of a related notion (though one that is subtly different, as discussed in [54]).

4.4 Properties of Protocols

Along with the growing number of dialogue protocols that have been suggested by various researchers comes the need to understand the properties of such protocols. Without this knowledge we have no basis for choosing between them, or even assessing whether they are adequate for a given purpose. Clearly it is possible, as in, to examine specific individual protocols and determine, for example, whether the dialogues that they enable will terminate [59], and what the possible outcomes of those dialogues are [58]. One severe difficulty with this, nevertheless, lies in the fact that it requires to make assumptions regarding agents’ attitudes towards the treatment of arguments, as detailed below. It is not the place here to enter into the details of such properties for specific interaction contexts, but we refer the reader to the recent survey by Henry Prakken on *persuasion dialogues* [65]. As for argument-based *negotiation*, we mention the very recent work by Amgoud and colleagues [2], which studies the properties of a (monotonic) bargaining protocol where agents only make concessions when they cannot defend their position any longer. This is an interesting attempt to formally extend the kind of results that are usually obtained in the context of such bilateral protocols (in particular regarding the optimality of the compromise) to a context where some sort of argumentation is permitted.

To conclude on these aspects, we mention two recent developments in this area that, we believe, pave the way for some potentially more foundational progresses in the near future.

- Firstly, the methodology adopted so far seems a rather unsatisfactory approach — it requires considerable theoretical work to be performed in order to understand any new protocol. Much more use would be to have a *meta-theory* of protocols which would identify the properties of a large class of protocols. Some tentative steps towards such a meta-theory are reported in [55].

---

⁴ [51] draws its title from the slogan, used in Britain during the Second World War, that “Loose lips sink ships” — a warning not to inadvertently give away information that might seem worthless but could prove fatal.

- Secondly, in order to assess the quality of (or bias induced by) a protocol, it is important to distinguish what is inherent to the problem itself; and what can really be imputed to the protocol. It is then very useful to be able to compare this protocol against an idealized situation where a fully-informed third-party would centrally compute the outcome (note that the bias induced may be interpreted as a quality loss, but can also sometimes be sought when viewed as a liberality offered to agents). Hence the need to be able to compute this centralized outcome. Sometimes this problem itself is challenging, for instance in a situation where several (potentially more than two) agents hold argumentation theories involving different sets of arguments and attack relations. A recent paper [16] explores this problem, and investigate the *merging* of several argumentation systems coming from different agents.

4.5 Argument Interchange Format

One major barrier to the development and practical deployment of argumentation systems is the lack of a shared, agreed notation for an “interchange format” for arguments and argumentation. Such a format is necessary if agents are to be able to exchange argumentative statements in open systems. The recently proposed Argument Interchange Format (AIF) [15] is intended to fill this gap, providing an approach to the representation and exchange of data between various argumentation tools and agent-based applications. It represents a consensus “abstract model” established by researchers across the fields of argumentation, artificial intelligence and multi-agent systems. The core AIF ontology is specified in a way that it can be extended to capture a variety of argumentation formalisms and schemes. One such extension, in the context of Semantic Web applications, deals with Walton’s theoretical model of argument schemes [72].

5 Concluding Remarks

Argumentation theory has been concerned with the study of rational human reasoning and dialogue for millennia. It is therefore an ideal resource for techniques, results and intuitions for problems in multi-agent reasoning and communication, and it is no surprise that formal models of argumentation are becoming an increasingly popular subject within research on multi-agent systems.

This chapter has presented a brief survey of a section of the work on argumentation in multi-agent systems, a section that encompasses the work that, in the opinion of the authors, is currently the most interesting of the work in the field. In short, in our view, the basic tools and methods have been established — we have well founded argumentation systems, and we have in dialogue games a means of structuring interactions between agents. What we need to do is to work with these tools in three directions. First, we need to integrate them into the reasoning processes of agents. For example, we need to decide how what an agent knows informs what it chooses to say in an interaction, and, conversely, what is

said in an interaction informs what an agent knows (some preliminary work on this later appeared in [57]). Second, we need to understand better how to design argumentation-based agent interactions so that they achieve the things that we want — we don't just need theoretical results that tell us how specific protocols work, but we need a theory that tells us how all protocols work. Third, we need to be able to show the effectiveness of argumentation-based agent interactions. In the end, however attractive the theory, if argumentation-based approaches are not more effective than other approaches to creating interactions between agents, then work on them is work wasted. As the paper surveyed above, and the work described in the contributions to this volume, show, as a community we are taking some steps in these three important directions.

Acknowledgements

The authors are grateful to the members of the Steering Committee and the Programme Committee of the International Workshop on Argumentation in Multi-Agent Systems (ArgMAS) for their support and advice. This survey paper is a revised and extended version of the editorial paper by Iyad Rahwan, that opens the Special Issue on Special Issue on Argumentation in Multi-Agent Systems of the Journal of Autonomous Agents and Multi-Agent Systems, Vol 11, No 2. (2005).

References

1. Amgoud, L.: A formal framework for handling conflicting desires. In: Nielsen, T.D., Zhang, N.L. (eds.) ECSQARU 2003. LNCS (LNAI), vol. 2711, pp. 552–563. Springer, Heidelberg (2003)
2. Amgoud, L., Dimopoulos, Y., Moraitis, P.: A unified and general framework for argumentation-based negotiation. In: AAMAS 2007, Honolulu, Hawaii, USA, ACM Press, New York (2007)
3. Amgoud, L., Kaci, S.: On the generation of bipolar goals in argumentation-based negotiation. In: Rahwan, I., Moraitis, P., Reed, C. (eds.) ArgMAS 2004. LNCS (LNAI), vol. 3366, Springer, Heidelberg (2005)
4. Amgoud, L., Maudet, N.: Strategical considerations for argumentative agents (preliminary report). In: Benferhat, S., Giunchiglia, E. (eds.) NMR 2002: Special session on Argument, Dialogue and Decision, pp. 399–407 (2002)
5. Amgoud, L., Maudet, N., Parsons, S.: Modelling dialogues using argumentation. In: Durfee, E. (ed.) ICMAS 1998, Boston MA, USA, pp. 31–38. IEEE Press, Los Alamitos (1998)
6. Amgoud, L., Parsons, S., Maudet, N.: Arguments, dialogue, and negotiation. In: Horn, W. (ed.) ECAI 2000, pp. 338–342. IOS Press, Amsterdam, Netherlands (2000)
7. Aristotle: Topics. In: Ross, W.D. (ed.) Clarendon, Oxford, UK (1928)
8. Atkinson, K., Bench-Capon, T., McBurney, P.: A dialogue game protocol for multi-agent argument over proposals for action. *Autonomous Agents and Multi-Agent Systems* 11(2), 153–171 (2006) Special issue on Argumentation in Multi-Agent Systems

9. Bench-Capon, T.J.M.: Argument in artificial intelligence and law. *Artificial Intelligence and Law* 5(4), 249–261 (1997)
10. Bourgne, G., Hette, G., Maudet, N., Pinson, S.: Hypothesis refinement under topological communication constraints. In: *AAMAS-2007*, Honolulu, Hawaii (May 2007)
11. Brewka, G.: *Nonmonotonic Reasoning: Logical Foundations of Commonsense*. Cambridge University Press, Cambridge, UK (1991)
12. Caminadas, M.: Semi-stable semantics. In: Dunne, P.E., Bench-Capon, T.J.M. (eds.) *COMMA 2006. Frontiers in Artificial Intelligence and Applications*, pp. 121–130. IOS Press, Amsterdam (2006)
13. Cayrol, C., Devred, C., Lagasque-Schiex, M.-C.: Handling controversial arguments in bipolar argumentation systems. In: Dunne, P.E., Bench-Capon, T.J.M. (eds.) *COMMA 2006. Frontiers in Artificial Intelligence and Applications*, pp. 261–272. IOS Press, Amsterdam (2006)
14. Chesñevar, C.I., Maguitman, A., Loui, R.P.: Logical models of argument. *ACM Computing Surveys* 32(4), 337–383 (2000)
15. Chesñevar, C.I., McGinnis, J., Modgil, S., Rahwan, I., Reed, C., Simari, G., South, M., Vreeswijk, G., Willmott, S.: Towards an argument interchange format. *The Knowledge Engineering Review* 21(4), 293–316 (2007)
16. Coste-Marquis, S., Devred, C., Konieczny, S., Lagasque-Schiex, M.-C., Marquis, P.: Merging argumentation systems. In: *AAAI 2005*, Pittsburgh, USA, pp. 614–619. AAAI Press, Stanford, California, USA (2005)
17. Dignum, F., Dunin-Kępicz, B., Berbrugge, R.: Agent theory for team formation by dialogue. In: Castelfranchi, C., Lespérance, Y. (eds.) *ATAL 2000. LNCS (LNAI)*, vol. 1986, pp. 150–166. Springer, Heidelberg (2001)
18. Dignum, F., Dunin-Kępicz, B., Berbrugge, R.: Creating collective intention through dialogue. *Logic Journal of the IGPL* 9(2), 289–303 (2001)
19. Doyle, J.: A truth maintenance system. *Artificial Intelligence* 12, 231–272 (1979)
20. Dung, P.M.: On the acceptability of arguments and its fundamental role in non-monotonic reasoning, logic programming and n-person games. *Artificial Intelligence* 77(2), 321–358 (1995)
21. Elhadad, M.: Using argumentation in text generation. *Journal of Pragmatics* 24, 189–220 (1995)
22. Elvang-Gøransson, M., Krause, P., Fox, J.: Acceptability of arguments as logical uncertainty. In: Moral, S., Kruse, R., Clarke, E. (eds.) *ECSQARU 1993. LNCS*, vol. 747, pp. 85–90. Springer, Heidelberg (1993)
23. Ferguson, G.: *Knowledge Representation and Reasoning for Mixed-Initiative Planning*. PhD thesis, Computer Science Department, University of Rochester, URCS TR 562 (January 1995)
24. Fernández, R., Endriss, U.: Abstract models for dialogue protocols. *Journal of Logic, Language and Information* 16(2), 121–140 (2007)
25. FIPA. *Communicative Act Library Specification*. Technical Report XC00037H, Foundation for Intelligent Physical Agents (August 10, 2001)
26. Fox, J., Barber, D., Bardhan, K.D.: Alternatives to Bayes? A quantitative comparison with rule-based diagnostic inference. *Methods of Information in Medicine* 19, 210–215 (1980)
27. Fox, J., Krause, P., Ambler, S.: Arguments, contradictions and practical reasoning. In: Neumann, B. (ed.) *ECAI-1992*, Vienna, Austria, pp. 623–627 (1992)
28. Fox, J., Parsons, S.: Arguing about beliefs and actions. In: Hunter, A., Parsons, S. (eds.) *Applications of Uncertainty Formalisms. LNCS (LNAI)*, vol. 1455, pp. 266–302. Springer, Heidelberg (1998)

29. Gärdenfors, P.: *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. MIT Press, Cambridge MA, USA (1988)
30. Georgeff, M.P.: Planning. *Annual Review of Computer Science* 2, 359–400 (1987)
31. Gordon, T.F., Karacapilidis, N.: The Zeno argumentation framework. In: *Proceedings of the Sixth International Conference on AI and Law*, pp. 10–18. ACM Press, New York (1997)
32. Hamblin, C.L.: *Fallacies*. Methuen, London, UK (1970)
33. Hitchcock, D., McBurney, P., Parsons, S.: A framework for deliberation dialogues. In: Hansen, H.V., Tindale, C.W., Blair, J.A., Johnson, R.H. (eds.) *OSSA 2001*, Ontario, Canada (2001)
34. Horvitz, E.J., Breese, J.S., Henrion, M.: Decision theory in expert systems and artificial intelligence. *International Journal of Approximate Reasoning* 2, 247–302 (1988)
35. Hulstijn, J.: *Dialogue models for enquiry and transaction*. PhD thesis, Universiteit Twente, Enschede, The Netherlands (2000)
36. Hulstijn, J., van der Torre, L.: Combining goal generation and planning in an argumentation framework. In: Hunter, A., Lang, J. (eds.) *NMR 2004*, Whistler, Canada (June 2004)
37. Jung, H., Tambe, M.: Towards argumentation as distributed constraint satisfaction. In: *Proceedings of AAAI Fall Symposium on Negotiation Methods for Autonomous Cooperative Systems*, AAAI Press, Stanford, California, USA (2001)
38. Jung, H., Tambe, M., Kulkarni, S.: Argumentation as distributed constraint satisfaction: applications and results. In: Müller, J.P., Andre, E., Sen, S., Frasson, C. (eds.) *Proceedings of the Fifth International Conference on Autonomous Agents*, Montreal, Canada, pp. 324–331. ACM Press, New York (2001)
39. Kakas, A.C., Moraitis, P.: Argumentation based decision making for autonomous agents. In: Rosenschein, J.S., Sandholm, T., Wooldridge, M., Yokoo, M. (eds.) *AAMAS-2003*, Melbourne, Victoria, pp. 883–890. ACM Press, New York (2003)
40. Karunatilake, N.C., Jennings, N.R.: Is it worth arguing? In: Rahwan, I., Moraitis, P., Reed, C. (eds.) *ArgMAS 2004*. LNCS (LNAI), vol. 3366, pp. 234–250. Springer, Heidelberg (2005)
41. Krause, P., Ambler, S., Elvang-Gøransson, M., Fox, J.: A logic of argumentation for reasoning under uncertainty. *Computational Intelligence* 11, 113–131 (1995)
42. Loui, R.P.: Defeat among arguments: a system of defeasible inference. *Computational Intelligence* 3, 100–106 (1987)
43. Loui, R.P.: Process and policy: Resource-bounded non-demonstrative reasoning. *Computational Intelligence* 14, 1–38 (1993)
44. Maudet, N., Chaib-draa, B.: Commitment-based and dialogue-game based protocols – new trends in agent communication language. *Knowledge Engineering Review* 17(2), 157–179 (2003)
45. Mbarki, M., Bentahar, J., Moulin, B.: Strategic and tactic reasoning for communicating agents. In: Maudet, N., Parsons, S., Rahwan, I. (eds.) *ArgMAS 2007*. LNCS (LNAI), vol. 4766, Springer, Heidelberg, Germany (2007)
46. McBurney, P., Parsons, S.: Dialogue game protocols. In: Huget, M.-P. (ed.) *Communication in Multiagent Systems*. LNCS (LNAI), vol. 2650, pp. 269–283. Springer, Heidelberg (2003)
47. McBurney, P., van Eijk, R.M., Parsons, S., Amgoud, L.: A dialogue-game protocol for agent purchase negotiations. *Journal of Autonomous Agents and Multi-Agent Systems* 7(3), 235–273 (2003)
48. McCarthy, J.: Circumscription – a form of non-monotonic reasoning. *Artificial Intelligence* 13, 27–39 (1980)

49. Nielsen, S.H., Parsons, S.: Computing preferred extensions for argumentation systems with sets of attacking arguments. In: Dunne, P.E., Bench-Capon, T.J.M. (eds.) *COMMA 2006*, pp. 97–108. IOS Press, Amsterdam (2006)
50. Ontañón, S., Plaza, E.: Arguments and counterexamples in case-based joint deliberation. In: Maudet, N., Parsons, S., Rahwan, I. (eds.) *ArgMAS 2007*. LNCS (LNAI), vol. 4766, Springer, Heidelberg, Germany (2007)
51. Oren, N., Norman, T.J., Preece, A.: Loose lips sink ships: A heuristic for argumentation. In: Maudet, N., Parsons, S., Rahwan, I. (eds.) *ArgMAS 2007*. LNCS (LNAI), vol. 4766, Springer, Heidelberg, Germany (2007)
52. Parsons, S., Fox, J.: Argumentation and decision making: A position paper. In: Gabbay, D.M., Ohlbach, H.J. (eds.) *FAPR 1996*. LNCS, vol. 1085, pp. 705–709. Springer, Heidelberg (1996)
53. Parsons, S., Green, S.: Argumentation and qualitative decision making. In: Hunter, A., Parsons, S. (eds.) *ECSQARU 1999*. LNCS (LNAI), vol. 1638, pp. 328–339. Springer, Heidelberg (1999)
54. Parsons, S., McBurney, P., Sklar, E., Wooldridge, M.: On the relevance of utterances in formal inter-agent dialogues. In: *AAMAS-2007*, Honolulu, HI (May 2007)
55. Parsons, S., McBurney, P., Wooldridge, M.: Some preliminary steps towards a meta-theory for formal inter-agent dialogues. In: Rahwan, I., Moraitis, P., Reed, C. (eds.) *ArgMAS 2004*. LNCS (LNAI), vol. 3366, Springer, Heidelberg (2005)
56. Parsons, S., Sierra, C., Jennings, N.: Agents that reason and negotiate by arguing. *Journal of Logic and Computation* 8(3), 261–292 (1998)
57. Parsons, S., Sklar, E.: How agents alter their beliefs after an argumentation-based dialogue. In: Parsons, S., Maudet, N., Moraitis, P., Rahwan, I. (eds.) *ArgMAS 2005*. LNCS (LNAI), vol. 4049, Springer, Heidelberg (2006)
58. Parsons, S., Wooldridge, M.J., Amgoud, L.: On the outcomes of formal inter-agent dialogues. In: Rosenschein, J., Sandholm, T., Wooldridge, M.J., Yokoo, M. (eds.) *AAMAS-2003*, pp. 616–623. ACM Press, New York (2003)
59. Parsons, S., Wooldridge, M.J., Amgoud, L.: Properties and complexity of formal inter-agent dialogues. *Journal of Logic and Computation* 13(3), 347–376 (2003)
60. Pollock, J.L.: Defeasible reasoning. *Cognitive Science* 11, 481–518 (1987)
61. Pollock, J.L.: The logical foundations of goal-regression planning in autonomous agents. *Artificial Intelligence* 106(2), 267–334 (1998)
62. Prakken, H.: On dialogue systems with speech acts, arguments, and counterarguments. In: Brewka, G., Moniz Pereira, L., Ojeda-Aciego, M., de Guzmán, I.P. (eds.) *JELIA 2000*. LNCS (LNAI), vol. 1919, Springer, Heidelberg (2000)
63. Prakken, H.: Relating protocols for dynamic dispute with logics for defeasible argumentation. *Synthese* 127, 187–219 (2001)
64. Prakken, H.: Coherence and flexibility in dialogue games for argumentation. *Journal of Logic and Computation* 15, 1009–1040 (2005)
65. Prakken, H.: Formal systems for persuasion dialogue. *The Knowledge Engineering Review* 21, 163–188 (2006)
66. Prakken, H., Reed, C., Walton, D.N.: Argumentation schemes and generalisations in reasoning about evidence. In: *Proceedings of the 9th international conference on artificial intelligence and law*, pp. 32–41. ACM Press, New York (2003)
67. Prakken, H., Sartor, G.: The role of logic in computational models of legal argument: a critical survey. In: Pauli, J. (ed.) *Learning-Based Robot Vision*. LNCS, vol. 2048, pp. 342–343. Springer, Heidelberg (2001)
68. Prakken, H., Vreeswijk, G.: Logics for defeasible argumentation. In: Gabbay, D., Guenther, F. (eds.) *Handbook of Philosophical Logic*, 2nd edn. vol. 4, pp. 219–318. Kluwer Academic Publishers, Dordrecht, Netherlands (2002)

69. Rahwan, I.: Interest-based Negotiation in Multi-Agent Systems. PhD thesis, Department of Information Systems, University of Melbourne, Melbourne, Australia (2004)
70. Rahwan, I., Pasquier, P., Sonenberg, L., Dignum, F.: On the benefits of exploiting underlying goals in argument-based negotiation. In: Holte, R.C., Howe, A. (eds.) *AAAI-2007*, Menlo Park CA, USA, AAAI Press, Stanford, California, USA (2007)
71. Rahwan, I., Ramchurn, S.D., Jennings, N.R., McBurney, P., Parsons, S., Sonenberg, L.: Argumentation based negotiation. *Knowledge Engineering Review* 18(4), 343–375 (2003)
72. Rahwan, I., Zablit, F., Reed, C.: Laying the foundations for a world wide argument web. *Artificial Intelligence* (to appear, 2007)
73. Rao, A.S., Georgeff, M.P.: BDI-agents: from theory to practice. In: *Proceedings of the First International Conference on Multiagent Systems*, San Francisco, USA (1995)
74. Reed, C., Norman, T.J.: *Argumentation Machines: New Frontiers in Argument and Computation*. Argumentation Library, vol. 9. Kluwer Academic Publishers, Dordrecht, The Netherlands (2004)
75. Reiter, R.: A logic for default reasoning. *Artificial Intelligence* 13, 81–132 (1980)
76. Rueda, S.V., García, A.J., Simari, G.R.: Argument-based negotiation among BDI agents. *Computer Science & Technology* 2(7) (2002)
77. Sadri, F., Toni, F., Torroni, P.: Logic agents, dialogues and negotiation: an abductive approach. In: Stathis, K., Schroeder, M. (eds.) *Proceedings of the AISB 2001 Symposium on Information Agents for E-Commerce* (2001)
78. Simari, G.R., Garcia, A.J., Capobianco, M.: Actions, planning and defeasible reasoning. In: *Proceedings of the 10th International Workshop on Non-Monotonic Reasoning*, Whistler BC, Canada, pp. 377–384 (2004)
79. Simari, G.R., Loui, R.P.: A mathematical treatment of defeasible reasoning and its implementation. *Artificial Intelligence* 53, 125–157 (1992)
80. Sycara, K.: The PERSUADER. In: Shapiro, D. (ed.) *The Encyclopedia of Artificial Intelligence*, John Wiley & Sons, Chichester (1992)
81. van Eemeren, F.H., Grootendorst, R.F., Henkemaans, F.S.: *Fundamentals of Argumentation Theory: A Handbook of Historical Backgrounds and Contemporary Applications*. Lawrence Erlbaum Associates, Hillsdale NJ, USA (1996)
82. Verheij, B.: Dialectical argumentation with argumentation schemes: An approach to legal logic. *Artificial Intelligence and Law* 11(1-2), 167–195 (2003)
83. Walton, D.N.: *Argumentation Schemes for Presumptive Reasoning*. Erlbaum, Mahwah, NJ, USA (1996)
84. Wurman, P.R., Wellman, M.P., Walsh, W.E.: A parametrization of the auction design space. *Games and Economic Behavior* 35(1-2), 304–338 (2001)