Proceedings of the International Multiconference on
Computer Science and Information Technology pp. 223–229

ISBN 978-83-60810-27-9
ISSN 1896-7094

Argumentative agents

(Invited Paper)

Francesca Toni
Department of Computing
Imperial College London, UK
Email: ft@imperial.ac.uk

***Abstract*—Argumentation, initially studied in philosophy and law, has been researched extensively in computing in the last decade, especially for inference, decision making and decision support, dialogue, and negotiation.**

This paper focuses on the use of argumentation to support intelligent agents in multi-agent systems, in general and in the ARGUGRID project[1] and Agreement Technology action[2]. In particular, the paper reviews how argumentation can help agents take decisions, either in isolation (by evaluating pros and cons of conflicting decisions) or in an open and dynamic environment (by assessing the validity of information they become aware of). It also illustrates how argumentation can support negotiation and conflict resolution amongst agents (by allowing them to exchange information and fill in gaps in their incomplete beliefs). Finally, the paper discusses how arguments can improve the assessment of the trustworthiness of agents in contract-regulated interactions (by supporting predictions on these agents' future behaviours).

I. Introduction

ARGUMENTATION, initially studied in philosophy and law, has been researched extensively in computing in the last decade, especially for inference, decision making and decision support, dialogue, and negotiation [1], [2], [3].

Simply stated, argumentation focuses on interactions where parties plead for and against some conclusion. In its most abstract form [4], an argumentation framework consists simply of a set of abstract *arguments* and a binary relation representing the *attacks* between the arguments. By instantiating the notion of arguments and the attack relation, different argument systems can be constructed, predominantly based upon logic. One such system is assumption-based argumentation (ABA) [5], [6]. Here, arguments are computed from a given set of *rules* and are supported by rules and *assumptions*. Also, an argument attacks another argument if the former supports a claim conflicting with some assumptions in the latter, where conflicts are given in terms of an underlying notion of *contrary* of assumptions. Rules, assumptions and contraries are defined in terms of an underlying logic *language*. Different choices for this language give different instances of ABA.

Argumentation provides a powerful mechanism for dealing with incomplete, possibly inconsistent information. It is also fundamental for the resolution of conflicts and differences of opinion amongst different parties. Further, it is useful for "explaining" outcomes generated automatically. As a consequence, argumentation is a useful mechanism to support several aspects of agents in multi-agent systems. Indeed, agents are goal-driven, self-contained entities with partial information on the environments in which they are situated (including other agents inhabiting these environments), with conflicting goals, but often in need to cooperate in order to achieve these goals (e.g. because resources controlled by other agents are needed to achieve these goals). Cooperation is supported by communication in multi-agent systems, and opinions as well as explanations are often exchanged amongst communicating agents.

The potential of *argumentative agents* has been/is being explored in two European activities: the EC-funded ARGUGRID project and the COST-funded Agreement Technologies action.

The *ARGUGRID project* developed a grid-based platform populated by rational decision-making agents associated with service requestors/providers and users [7]. Within agents, argumentation as envisaged in ABA is used to support decision making, taking into account (and despite) the often conflicting information that these agents have, as well as the preferences of users, service requestors and providers [8], [9], [10]. In the ARGUGRID platform, argumentation is also used to support the negotiation between agents [10], [11] on behalf of service requestors/providers/users. This negotiation takes place within dynamically formed virtual organisations [12]. The agreed combination of services, amongst the argumentative agents, can be seen as a complex service within a service-oriented architecture [13]. The ARGUGRID approach has been validated by way of industrial application scenarios in e-procurement and earth observation [7], [8], [9].

The *Agreement Technologies action* aims at developing computer systems in which autonomous agents negotiate with one another, typically on behalf of humans, in order to come to mutually acceptable agreements [14], [15]. Agreement Technologies include argumentation, negotiation, and trust computing, as well as combinations of these.

In this paper we review some of the achievement to date of these activities, in providing argumentative agents and validating them against other approaches and in applications.

The paper is structured as follows. In section II we give some background on abstract argumentation and ABA. In section III we illustrate ways in which argumentative agents can take decisions. In section IV and V we describe the use of ABA for conflict resolution and negotiation, respectively, amongst argumentative agents. In section VI we review a

[1] www.argugrid.eu
[2] www.agreement-technologies.eu

possible integration of arguments with statistical information in trust computing. In section VII we conclude.

II. ARGUMENTATION

This section gives essential background on logic-based argumentation, focusing on abstract argumentation [4] and assumption-based argumentation [5], [6].

Abstract argumentation frameworks [4] are pairs (Arg, att) where Arg is a set of arguments and $att \subseteq Arg \times Arg$ is a binary relation representing attacks between arguments.

The main purpose of argumentation theory is to identify which arguments in an argumentation framework are rationally "acceptable". Several notions of acceptability have been proposed in the literature on argumentation, some providing an "intrinsic" measure of argument strength (e.g. [16]), whereby the acceptability of an argument depends on its internal logical structure, others giving a "dialectical" measure (e.g. [4], [17], [18], [19]), depending exclusively on attacking arguments.

An example of dialectical measure for abstract argumentation frameworks is given by *conflict-free extensions*, namely sets X of arguments such that there is no argument in X which attacks another argument in X. Another example is given by *admissible extensions* [4], namely sets X of arguments that are conflict-free and capable of defending themselves against every attacking argument (namely for every argument Y that attacks X, there is some argument in X that attacks Y). A further example is *preferred extensions* [4], namely (subset) maximal admissible extensions. These examples of dialectical measures are all "qualitative", based predominantly on the capability of arguments to defend themselves. "Quantitative" dialectical measures have been proposed too (e.g. [19]).

Assumption-based argumentation (ABA) frameworks [5], [6] can be defined for any logic specified by means of (inference) *rules*, by identifying sentences in the underlying logic *language* that can be treated as *assumptions*. Intuitively, *arguments* are deductions (in the chosen logic language) of a conclusion (or claim) supported by a set of assumptions. Then, *attacks* against arguments are always directed at the assumptions supporting the arguments. More precisely, an attack by one argument against another is a deduction by the first argument of the *contrary* of an assumption supporting the second argument.

The inference rules may be domain-specific or domain-independent, and may represent, for example, causal information, or laws and regulations. Assumptions are sentences in the language that are open to challenge, for example uncertain beliefs ("it will rain"), unsupported beliefs ("I believe that some service provider is reliable"), or decisions ("I will purchase a specific service"). Typically, assumptions can occur as premises of inference rules, but not as conclusions. In general, the contrary of an assumption is a sentence representing a challenge against the assumption. For example, the contrary of the assumption "it will rain" might be "the sky is clear". The contrary of the assumption 'I will purchase a specific service" might be "I will purchase a different service" (where I only need one service). The contrary of the assumption "I believe that some service provider is reliable" might be "there is evidence against that service provider being reliable".

Given arguments and attacks, several qualitative dialectical measures of acceptability have been deployed within ABA [5], [17], [6], including conflict-free, admissible and preferred extensions. Query answering with respect to these dialectical measures is implemented, for any ABA framework given as input, in the CaSAPI system[3] [20], [21], [22]. Here, queries represent claims whose dialectical validity with respect to a chosen notion of extension is under scrutiny.

III. ARGUMENTATION FOR DECISION-MAKING

Qualitative decision theory [23] has been advocated for quite some time as a viable and useful alternative to classical quantitative decision theory [24], when a decision problem cannot be easily formulated in standard decision-theoretic terms using decision tables, utility functions and probability distributions. A number of qualitative approaches to decision making, e.g. [25], [26], [27], [28], have been put forward, with argumentation-based decision making amongst them (e.g. [27]). In decision-theoretic terms, argumentation can be used as a model to compute a utility function which is too complex to be given a simple analytical expression in closed form. Argumentation has been proposed for decision making under certainty (where the outcomes of decisions are known to the decision maker) [8], [9], strict uncertainty (where the outcomes of decisions are uncertain and no probabilistic information is available) [29], [30], [31], [32], [10], and also for decision under risk (where some probabilistic information is known) [16], [33], [34]. Argumentation has also been used to support practical reasoning [35], [36], and decision support systems [37], [38], [39]. Further, arguments can be seen as supporting "values", as in value-based argumentation for decision-making [40]. Finally, argumentation can be used for computing decision tables, utility functions and probability distributions in classical quantitative decision theory [41].

Here we summarise two different uses of ABA (see section II) to support decision making under certainty [8], [9] and under strict uncertainty [10].

We consider the following decision problems:

- Let $\mathcal{D}$ be a (non-empty) set of alternative decisions.
- The outcomes of decisions are individual states $s \in \mathcal{S}$ (if under certainty) or sets of states $S \subseteq \mathcal{S}$ (if under uncertainty).
- States can be seen as sets of "goals" of (or benefits for) the decision maker, which can be represented as sentences in a given set $\mathcal{G}$.
- Preferences over goals may be optionally specified, e.g. in the form of weights (positive integers). These can be expressed by a mapping $w : \mathcal{G} \to \mathbb{N}$. The case where all weights are the same is equivalent to the case where no weights are specified.

[3] http://www.doc.ic.ac.uk/~dg00/casapi.html

- Rather than being known a-priori, outcomes of decisions are determined from a belief base $\mathcal{B}$. The beliefs "entailed" by this base correspond to states.

Decisions may correspond to products, including e-procurement products [8], earth observation products [9], commodities [10] etc. For these decision problems, argumentation can be used to determine the relative value of different decisions, to single out decisions with "top-most" value and for explaining decisions.

Under certainty, and assuming all goals have equal weight, decisions with top-most value can be defined as "dominant" decisions, where

- a decision $d \in \mathcal{D}$ is *dominant* if and only if the outcome $s \in \mathcal{S}$ of d is such that, for any alternative decision $d' \in \mathcal{D}$, if $s' \in \mathcal{S}$ is the outcome of d', then $s \supseteq s'$.

This is the approach taken in [9].

Alternatively, decisions with top-most value are decisions resulting in states that are upper bounds of partial orders over states. Under certainty, a partial order $\sqsupseteq$ over states can be given as follows:

- a state $s \in \mathcal{S}$ is *strictly preferred to* a state $s' \in \mathcal{S}$ (denoted $s \sqsupset s'$) if and only if
  1) there exists a goal $g \in \mathcal{G}$ such that $g \in s$ but $g \notin s'$, and
  2) for each goal $g' \in \mathcal{G}$, if $w(g') \geq w(g)$ and $g' \in s'$ then $g' \in s$;
- a state $s \in \mathcal{S}$ is *preferred to* a state $s' \in \mathcal{S}$ (denoted $s \sqsupseteq s'$) if and only if $s \sqsupset s'$ or $s = s'$.

This partial order is used in [10] to define a partial order over decisions under strict uncertainty, as we will see below.

Note that, when all weights are the same, $s \sqsupseteq s'$ is equivalent to $s \supseteq s'$.

In [8], [9], the belief base maps features of products to goals of the decision maker. For example, a hotel with rooms costing less than 50£ may be believed to be cheap (where the price is a feature and being cheap is a goal) [9]. Further, in the case of e-procurement for an e-ordering system [8], an e-ordering system with a 3-year flat cost may be deemed to decrease costs. Here, features determine univocally (with certainty) goals. The belief base is represented as an ABA framework from which the following arguments can be built:

(i) "choose decision d because d allows to achieve goal g"
(ii) "do not choose d because some other decision allows to achieve goal g and I am not sure d does"

(see [8], [9] for formal details). Arguments of type (ii) attack arguments of type (i) and vice versa. Then, dominant decisions, as given above, correspond to admissible sets of arguments for the given ABA frameworks. ABA thus allows to compute dominant decisions and explain these decisions (by presenting the arguments). Moreover, in [9] we also propose a different argumentation semantics based on counting (and resulting in a numerical, rather than Boolean value for arguments) and links this semantics to dominant decisions when weights are given (see [9] for details).

In [10], the belief base encodes again a mapping between features and goals, but using information that may be incomplete (e.g. that a good school is located in the vicinity of a real estate property) and that may lead to conflicts/inconsistencies (e.g. that a real estate property is in a safe area and is not in a safe area).[4] As a consequence, decisions correspond to sets of states, where different states correspond to different assumptions (completing the information) and different resolutions of the conflicts. These resolutions (states) are preferred extensions, in the argumentation sense (see section II), and they can be compared using the standard minimax criterion from decision theory, using the following notations:

- for a given decision $d \in \mathcal{D}$, let $cred\_pref(d)$ be the set of all $s \in \mathcal{S}$ such that s is satisfied in some preferred extension of the belief base extended by d;
- for any set of states $S \subseteq \mathcal{S}$, let $min(S)$ be a state such that for each goal $g \in \mathcal{G}$, g is satisfied in $min(S)$ if and only if g is satisfied in every state in S.

Then

- a decision $d \in \mathcal{D}$ is *minimax-preferred to* a decision $d' \in \mathcal{D}$ if and only if
  – $min(cred\_pref(d)) \sqsupseteq min(cred\_pref(d'))$

  where $\sqsupseteq$ is as defined earlier.

This notion of minimax-preference is equivalent to a purely argumentative preference notion between sets of states, defined using the following notion:

- for a given decision $d \in \mathcal{D}$, let $scept\_pref\_state(d)$ be the state $s \in \mathcal{S}$ consisting of all goals holding in all preferred extensions of the belief base extended by d.

Then,

- a decision $d \in \mathcal{D}$ is *sceptically-preferred to* a decision $d' \in \mathcal{D}$ if and only if
  – $scept\_pref(d) \sqsupseteq scept\_pref(d')$.

The fully argumentative notion of sceptically-preferred and the partially argumentative, partially decision-theoretic notion of minimax-preferred are equivalent [10]. Top-elements of the partial orders given by either notions are decisions with top-most values.

Overall, the two approaches considered use argumentation for different purposes: on one side, [8], [9] encodes a fully decision-theoretic notion of dominance into an ABA framework, and uses argumentation to "explain" dominant decisions under certainty; on the other side, [10] uses argumentation to deal with conflicts and incomplete information for decisions under strict uncertainty, and a sceptical semantics to mirror a minimax decision-theoretic criterion.

IV. Argumentation for conflict resolution

Complex multi-agent systems are composed of heterogeneous agents with different, possibly incomplete beliefs and

[4] Note that, in ABA, assumptions are used as premises of rules to represent incomplete information that can be "completed" by making suitable assumptions. Also, assumptions are used to render rules defeasible, thus paving the way to resolving inconsistencies.

different, possibly conflicting desires. Conflicts may thus arise amongst agents for (at least) two reasons. Firstly, agents may make different assumptions to fill gaps in their beliefs, where some of these assumptions may be incorrect, and decide on incompatible, conflicting actions due to misinformation. Secondly, even if agents share the same information, they may still disagree if they have conflicting desires.

Due to ABA's suitability in dealing with incomplete and conflicting information, agents' beliefs and desires can be represented in ABA [42], [43]. Following an existing trend of work in argumentation for conflict resolution (e.g. see [44]), in [45] we use ABA to resolve conflicts between *two agents*. These conflicts arise when the agents have different goals, g_1 and g_2, and different decisions, d_1 and d_2, having those goals as respective outcomes, according to their respective individual belief bases. These bases are assumed to be represented as ABA frameworks. Here, rules are used to represent beliefs about the achievement of desires as well as factual information.

For both agents, the goal belongs to a conflict-free extension of the agent's ABA framework, extended with the agent's decision. The agents' objective is to resolve the conflict, by agreeing on a common goal g and a common decision d such that g is a variant of both g_1 and g_2. For example, in a service-oriented architecture, if the two agents represent two service requestors from the same organisation, with different requirements but with the shared goal of obtaining some service of a certain type, then

- g_1 may correspond to obtaining a service s_1 of that type, by purchasing it (decision d_1),
- g_2 may correspond to obtaining another service s_2 of the same type, by purchasing it (decision d_2), and
- g may correspond to obtaining a service s of that type, by purchasing it (decision d), where s may be one of s_1 or s_2 or a new service.

Here, the original choice of s_1/s_2 by the agents may be dictated by their lack of knowledge of the other agents requirements or of the availability and characteristics of services. For example, the first agent may know that s_1 fulfils some requirement of the second agent while the latter may be unaware of this. By passing information to the second agent, the first agent may be able to persuade the second to purchase s_1 (in this case s would be s_1). Thus, conflict resolution amounts to identifying a goal that is the outcome of a decision in conflict-free extensions of either belief bases (ABA frameworks) after the agents have shared factual information (e.g. that s_1 fulfils some requirement).

Alternatively, this conflict resolution can be achieved by "merging" the two ABA frameworks. The merge eliminates misunderstanding between agents, allows to revise the agents' incorrect assumptions and takes into account desires from both agents. To satisfy desires from both agents, the mechanism of concatenation is used to merge rules. Upon a successful concatenation merge, both agents' desires may be satisfied (if they can be satisfied). Details of this approach can be found in [45]. Here, a dialogical counterpart to the merge is also sketched as a more realistic approach to conflict resolution.

V. Argumentation for negotiation

The need for negotiation arises when autonomous agents have conflicting interests/desires but may benefit from cooperation in order to achieve them. In particular, this cooperation may amount to a change of goals (as in conflict-resolution, see section IV) and/or to the introduction of new goals (e.g. for an agent to provide a certain resource to another, even though it may not have originally planned to do so). Typically negotiation involves (fair) compromise.

Argumentation-based negotiation is a particular class of negotiation, whereby agents can provide arguments and justifications as part of the negotiation process [46]. It is widely believed that the use of argumentation during negotiation increases the likelihood and/or speed of agreements being reached [47]. Argumentation can be used to support the decision-making taking place prior to or during negotiation. Moreover, argumentation can be used to conduct negotiation, by supporting the resolution of conflicts giving rise to the need of negotiation and by filling in information gaps and rectifying misinformed beliefs.

In [10] we propose the use of ABA to support decision making under strict uncertainty (as described in section III) prior to the agents engaging in negotiation. The negotiation takes place between a buyer and a seller (e.g. of services) and results in (specific forms of) contracts, taking into account contractual properties and preferences that buyer and seller have. The negotiation is guided by a "minimal concession" strategy that is proven to be in symmetric Nash equilibrium. Adopting this strategy, agents may concede and adopt a less-preferred goal to the one they currently hold (namely a goal with a smaller weight, according to the presentation in section III) for the sake of reaching agreement. Thus, agreement amounts to compromise. This approach has been extended in [48] to incorporate rewards during negotiation. These rewards in turn can be seen as arguments in favour of agreement.

In [11] we study the use of a form of ABA, given in [49], for improved effectiveness of the negotiation process, in particular concerning the number of dialogues and dialogue moves that need to be performed during negotiation without affecting the quality of solutions reached. The focus here is on negotiation of resources in resource reallocation settings. This work complements studies on protocols for argumentation-based negotiation (e.g. [50]) and argumentation-based decision making during negotiation (e.g. [10]) by integrating argumentation-based decision making with the exchange of arguments to benefit the outcome of negotiation. Agents engage in dialogues with other agents in order to obtain resources they need but do not have. Dialogues are regulated by simple communication policies that allow agents to provide reasons (arguments) for their refusals to give away resources; agents use ABA in order to deploy these policies. We assess the benefits of providing these reasons both informally and experimentally: by providing reasons, agents are more effective in identifying a reallocation of resources if one exists and failing if none exists.

VI. Argumentation for trust computing

Computing trust is a problem of reasoning under uncertainty, requiring the prediction and anticipation by an agent (the evaluator) of the future behaviour of another agent (the target). Despite the acknowledged ability of argumentation to support reasoning under uncertainty (e.g. see [16]), only Prade [51], Dondio & Barret [52] and Parsons et al [53] have considered the use of arguments for computing trust in a local trust rating setting. Dondio & Barret [52] propose a set of trust schemes, in the spirit of Walton's argument schemes [54], and assume a dialectical process between the evaluator and the target whereby the evaluator poses critical questions against arguments by the target concerning its trustworthiness. Prade [51] proposes an argumentation-based approach for trust evaluation that is bipolar (separating arguments for trust and for distrust) and qualitative (as arguments can support various degrees of trust/distrust). Parsons et al [53] define an argumentation logic where arguments support measures of trust, e.g. qualitative measures such as "very reliable" or "somewhat unreliable".

There are several non-argumentation based methods to model the trust of the evaluator in the target. Sabater and Sierra [55] classify approaches to trust as either "cognitive", based on underlying beliefs, or "game-theoretical", where trust values correspond to subjective probabilities and can be modelled by uncertainty values, Bayesian probabilities, fuzzy sets, or Dempster-Shafer belief functions. The latter approach is predominant nowadays for trust computing. However, Castelfranchi and Falcone [56] argue against a purely game-theoretic approach to trust and in favour of a cognitive approach based upon a mental model of the evaluator, including goals and beliefs. Moreover, some works (e.g. [57]) advocate the need for and benefits of hybrid trust models, combining both the cognitive and game-theoretical approach.

In recent work [58], we propose a hybrid approach for constructing Dempster-Shafer belief functions modeling the trust of the evaluator in the target by combining statistical information concerning the past behaviour of the target and arguments concerning the target's expected behaviour. These arguments are built from current and past contracts between evaluator and target, and are integrated with statistical information proportionally to their validity. Concretely, in a service-oriented setting, the statistics are drawn from past behaviour of the target in the delivery of agreed services and, following [59], a classification of this behaviour as "good" (the service was delivered as agreed), "bad" (the service was not delivered as agreed) or "inappreciable" (the evaluator cannot judge the delivery of the service). Clearly, the more "good" behaviour the target has shown in the past the more likely the evaluator will be to trust it. The arguments are drawn from contracts regulating the delivery of services, as follows:

- a forecast argument supporting the claim of not trusting the target (as far as delivering a service is concerned) if there is no guarantee on the quality of that service in the form of a written contract clause;
- a forecast argument supporting the claim of trusting the target if there exists a guarantee in the form of a contract clause;
- an argument attacking the forecast argument for trusting the target if the target has in the past "most often" violated existing contract clauses.

The applicable arguments and attacks form an abstract argumentation framework (see section II). They are combined with the statistics in accordance to their strength, measured using the method of [19].

This method of measuring trust extends a standard method for trust [59] that relies upon the statistical information only. The two methods have identical predictive performance when the evaluator is highly "cautious", but the hybrid method gives a significant increase when the evaluator is not or is only moderately "cautious". Moreover, with the hybrid method, target agents are more motivated to honour contracts than when trust is computed on a purely statistical basis. The comparison between the two methods is performed within a simulated setting (see [58] for details).

VII. Conclusion

Argumentation, initially studied in philosophy and law, has been researched extensively in computing in the last decade, especially for inference, decision making and decision support, dialogue, and negotiation.

This paper has summarised some of the uses of argumentation

(i) to help agents to make decision, either in isolation (by evaluating pros and cons of conflicting decisions) or in an open and dynamic environment (by assessing the validity of information they become aware of)
(ii) to support negotiation and
(iii) conflict resolution amongst agents, and
(iv) to improve the assessment of the trustworthiness of agents in contract-regulated interactions.

The paper has focused on contributions to (i)–(iv) developed within two European initiatives: the EC-funded ARGUGRID project and the COST-funded Agreement Technologies action.

Acknowledgment

The author would like to thank the Agreement Technology COST action IC 0801 for partially sponsoring this work.

References

[1] C. I. Chesñevar, A. G. Maguitman, and R. P. Loui, "Logical models of argument," *ACM Computing Surveys*, vol. 32, no. 4, pp. 337–383, 2000.
[2] T. Bench-Capon and P. E. Dunne, "Argumentation in artificial intelligence," *Artificial Intelligence*, no. 171, pp. 619–641, 2007.

[3] I. Rahwan and P. McBurney, "Guest editors' introduction: Argumentation technology," *IEEE Intelligent Systems*, vol. 22, no. 6, pp. 21–23, 2007.

[4] P. Dung, "On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming, and n-person games," *Artificial Intelligence*, vol. 77, no. 2, pp. 321–257, 1995.

[5] A. Bondarenko, P. M. Dung, R. A. Kowalski, and F. Toni, "An abstract, argumentation-theoretic approach to default reasoning," *Artificial Intelligence*, vol. 93, no. 1–2, pp. 63–101, 1997.

[6] P. Dung, R. Kowalski, and F. Toni, "Assumption-based argumentation," in *Argumentation in AI*, I. Rahwan and G. Simari, Eds. Springer-Verlag, 2009, pp. 199–218.

[7] F. Toni, M. Grammatikou, S. Kafetzoglou, L. Lymberopoulos, S. Papavassileiou, D. Gaertner, M. Morge, S. Bromuri, J. McGinnis, K. Stathis, V. Curcin, M. Ghanem, and L. Guo, "The ArguGRID platform: An overview," in *Proceedings of Grid Economics and Business Models, 5th International Workshop (GECON 2008)*, ser. Lecture Notes in Computer Science, J. Altmann, D. Neumann, and T. Fahringer, Eds., vol. 5206. Springer, August 2008, pp. 217–225.

[8] P.-A. Matt, F. Toni, T. Stournaras, and D. Dimitrelos, "Argumentation-based agents for eprocurement," in *Proceedings of the 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)- Industry and Applications Track*, M. Berger, B. Burg, and S. Nishiyama, Eds., 2008, pp. 71–74.

[9] P.-A. Matt, F. Toni, and J. Vaccari, "Dominant decisions by argumentation agents," in *Proceedings of the Sixth International Workshop on Argumentation in Multi-Agent Systems (ArgMAS 2009), affiliated to AAMAS 2009*, ser. Lecture Notes in Computer Science, P. McBurney, I. Rahwan, S. Parsons, and N. Maudet, Eds., vol. 6057. Springer, 2010, pp. 42–59.

[10] P. M. Dung, P. M. Thang, and F. Toni, "Towards argumentation-based contract negotiation," in *Proceedings of the 2nd International Conference on Computational Models of Argument (COMMA'08)*, ser. Frontiers in Artificial Intelligence and Applications, P. Besnard, S. Doutre, and A. Hunter, Eds., vol. 172. IOS Press, 2008, pp. 134–146.

[11] A. Hussain and F. Toni, "On the benefits of argumentation for negotiation - preliminary version," in *Proceedings of 6th European Workshop on Multi-Agent Systems (EUMAS-2008)*, 2008.

[12] J. McGinnis, K. Stathis, and F. Toni, "A formal framework of virtual organisations as agent societies," *EPTCS*, vol. 16, p. 1, 2010.

[13] F. Toni, "Argumentative KGP agents for service composition," in *Proc. AITA08, Architectures for Intelligent Theory-Based Agents, AAAI Spring Symposium*, M. Balduccini and C. Baral, Eds. Stanford University, 2008.

[14] S. Ossowski, "Coordination and agreement in multi-agent systems," in *Proceedings of Cooperative Information Agents XII, 12th International Workshop (CIA 2008)*, ser. Lecture Notes in Computer Science, M. Klusch, M. Pechoucek, and A. Polleres, Eds., vol. 5180. Springer, 2008, pp. 16–23.

[15] ——, "Coordination in multi-agent systems: Towards a technology of agreement," in *Proceedings of Multiagent System Technologies, 6th German Conference (MATES 2008)*, ser. Lecture Notes in Computer Science, R. Bergmann, G. Lindemann, S. Kirn, and M. Pechoucek, Eds., vol. 5244. Springer, 2008, pp. 2–12.

[16] P. Krause, S. Ambler, M. Elvang-Gøransson, and J. Fox, "A logic of argumentation for reasoning under uncertainty," *Computational Intelligence*, vol. 11, pp. 113–131, 1995.

[17] P. Dung, P. Mancarella, and F. Toni, "Computing ideal sceptical argumentation," *Artificial Intelligence*, vol. 171, no. 10–15, pp. 642–674, 2007.

[18] M. Caminada, "Semi-stable semantics," in *Proceedings of 1st International Conference on Computational Models of Argument (COMMA'06)*, ser. Frontiers in Artificial Intelligence and Applications, P. E. Dunne and T. J. M. Bench-Capon, Eds., vol. 144. IOS Press, 2006, pp. 121–130.

[19] P.-A. Matt and F. Toni, "A game-theoretic measure of argument strength for abstract argumentation," in *Proceedings of 11th European Conference on Logics in Artificial Intelligence (JELIA 2008)*, ser. Lecture Notes in Computer Science, S. Hölldobler, C. Lutz, and H. Wansing, Eds., vol. 5293, 2008, pp. 285–297.

[20] D. Gaertner and F. Toni, "CaSAPI - a system for credulous and sceptical argumentation," in *Proceedings of the International Workshop on Argumentation and Non-Monotonic Reasoning (ArgNMR 2007), affiliated to LPNMR 2007*, G. Simari and P. Torroni, Eds., 2007.

[21] ——, "On computing arguments and attacks in assumption-based argumentation," *IEEE Intelligent Systems, Special Issue on Argumentation Technology*, vol. 22, no. 6, pp. 24–33, November/December 2007.

[22] ——, "Hybrid argumentation and its properties," in *Proceedings of the Second International Conference on Computational Models of Argument (COMMA'08)*, ser. Frontiers in Artificial Intelligence and Applications, P. Besnard, S. Doutre, and A. Hunter, Eds., vol. 172. IOS Press, 2008, pp. 183–195.

[23] J. Doyle and R. H. Thomason, "Background to qualitative decision theory," *AI Magazine*, vol. 20, no. 2, pp. 55–68, 1999.

[24] S. French, *Decision theory: an introduction to the mathematics of rationality*. Ellis Horwood, 1987.

[25] J. Pearl, "From conditional oughts to qualitative decision theory," in *Proceedings of the 9th Conference on Uncertainty in Artificial Intelligence (UAI'93)*, D. Heckerman and E. H. Mamdani, Eds. Morgan Kaufmann, 1993, pp. 12–22.

[26] D. Poole, "Probabilistic Horn abduction and Bayesian networks," *Artificial Intelligence*, vol. 64, no. 1, pp. 81–129, 1993.

[27] B. Bonet and H. Geffner, "Arguing for decisions: A qualitative model of decision making," in *Proceedings of the 12th Conference on Uncertainty in Artificial Intelligence (UAI-96)*, E. Horvitz and F. V. Jensen, Eds. Morgan Kaufmann, 1996, pp. 98–105.

[28] D. Dubois, H. Fargier, and P. Perny, "Qualitative decision theory with preference relations and comparative uncertainty: An axiomatic approach," *Artificial Intelligence*, vol. 148, pp. 219–260, 2003.

[29] J. Fox, P. Krause, and M. Elvang-Gøransson, "Argumentation as a general framework for uncertain reasoning," in *Proceedings of the 9th Conference on Uncertainty in Artificial Intelligence (UAI'93)*, D. Heckerman and E. H. Mamdani, Eds. Morgan Kaufmann, 1993, pp. 428–434.

[30] J. Fox and S. Parsons, "On using arguments for reasoning about actions and values," in *Working Papers of the AAAI Spring Symposium on Qualitative Preferences in Deliberation and Practical Reasoning*, J. Doyle and R. H. Thomason, Eds., 1997, pp. 55–63.

[31] L. Amgoud, "A unified setting for inference and decision: An argumentation-based approach," in *proceedings of the 21st Conference on Uncertainty in Artificial Intelligence (UAI'05)*. AUAI Press, 2005, pp. 26–33.

[32] L. Amgoud and H. Prade, "Making decisions from weighted arguments," in *Decision theory and multi-agent planning*, G. D. Riccia, D. Dubois, R. Kruse, and H.-J. Lenz, Eds. Springer, 2006, pp. 1–14.

[33] S. Parsons, "Normative argumentation and qualitative probability," in *Proceedings of the First International Joint Conference on Qualitative and Quantitative Practical Reasoning (ECSQARU-FAPR'97)*, ser. Lecture Notes in AI, D. M. Gabbay, R. Kruse, A. Nonnengart, and H. J. Ohlbach, Eds., vol. 1244. Springer, jun 9–12 1997, pp. 466–480.

[34] L. Amgoud and H. Prade, "Using arguments for making decisions: A possibilistic logic approach," in *Proceedings of the 20th Conference of Uncertainty in Artificial Intelligence (UAI'04)*, D. M. Chickering and J. Y. Halpern, Eds. AUAI Press, 2004, pp. 10–17.

[35] H. Prakken, "Combining sceptical epistemic reasoning with credulous practical reasoning," in *Proceedings of 1st International Conference on Computational Models of Argument (COMMA'06)*, ser. Frontiers in Artificial Intelligence and Applications, P. Dunne and T. J. M. Bench-Capon, Eds., vol. 144. IOS Press, 2006, pp. 311–322.

[36] I. Rahwan and L. Amgoud, "An argumentation-based approach for practical reasoning," in *Proceedings of the 5th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'06)*, H. Nakashima, M. P. Wellman, G. Weiss, and P. Stone, Eds. ACM, 2006, pp. 347–354.

[37] J. Fox, N. Johns, C. Lyons, A. Rahmanzadeh, R. Thomson, and P. Wilson, "PROforma: a general technology for clinical decision support systems," *Computer Methods and Programs in Biomedicine*, vol. 54, no. 10–15, pp. 59–67, 1997.

[38] S. Modgil and P. Hammond, "Decision support tools for clinical trial design," *Artificial Intelligence in Medicine*, vol. 27, no. 2, pp. 181–200, 2003.

[39] M. Morge and P. Mancarella, "The hedgehog and the fox. An argumentation-based decision support system," in *Proceedings of the 4th International Workshop on Argumentation in Multi-Agent Systems (ArgMAS'07), affiliated to AAMAS*, ser. Lecture Notes in Computer Science, I. Rahwan, S. Parsons, and C. Reed, Eds., vol. 4946. Springer, 2008, pp. 114–131.

[40] F. S. Nawwab, T. J. M. Bench-Capon, and P. E. Dunne, "A methodology for action-selection using value-based argumentation," in *Proceedings of the Second International Conference on Computational Models of Argument (COMMA'08)*, ser. Frontiers in Artificial Intelligence and

Applications, P. Besnard, S. Doutre, and A. Hunter, Eds., vol. 172. IOS Press, 2008, pp. 264–275.

[41] P.-A. Matt, "Argumentation as a practical foundation for decision theory," Ph.D. dissertation, Department of Computing, Imperial College London, 2010.

[42] F. Toni, "Assumption-based argumentation for selection and composition of services," in *Proceedings of the 8th International Workshop on Computational Logic in Multi-Agent Systems (CLIMA VIII)*, ser. Lecture Notes in AI, F. Sadri and K. Satoh, Eds., vol. 5056. Springer, 2008, pp. 231–247.

[43] ——, "Assumption-based argumentation for closed and consistent defeasible reasoning," in *Proceedings of the First International Workshop on Juris-informatics (JURISIN 2007), in association with The 21th Annual Conference of The Japanese Society for Artificial Intelligence (JSAI2007)*, ser. Lecture Notes in Computer Science, K. Satoh, A. Inokuchi, K. Nagao, and T. Kawamura, Eds., vol. 4914. Springer, 2008, pp. 390–402.

[44] L. Amgoud and S. Kaci, "An argumentation framework for merging conflicting knowledge bases," *International Journal on Approximate Reasoning*, vol. 45, no. 2, pp. 321–340, 2007.

[45] X. Fan, F. Toni, and A. Hussain, "Two-agent conflict resolution with assumption-based argumentation," in *Proceedings of the Third International Conference on Computational Models of Argument (COMMA'10)*, P. Baroni, F. Cerutti, M. Giacomin, and G. Simari, Eds. IOS Press, 2010, pp. 231–242.

[46] N. R. Jennings, P. Faratin, A. R. Lomuscio, S. Parsons, C. Sierra, and M. Wooldridge, "Automated negotiation: Prospects, methods and challenges," *Group Decision and Negotiation*, vol. 10, no. 2, pp. 199–215, 2001.

[47] I. Rahwan, S. D. Ramchurn, N. R. Jennings, P. McBurney, S. Parsons, and L. Sonenberg, "Argumentation-based negotiation," *The Knowledge Engineering Review*, vol. 18, no. 4, pp. 343–375, 2004.

[48] P. M. Dung, P. M. Thang, and N. D. Hung, "Argument-based decision making and negotiation in e-business: Contracting a land lease for a computer assembly plant," in *Proceedings of the 9th International Workshop on Computational Logic in Multi-Agent Systems (CLIMA IX)*, ser. Lecture Notes in Computer Science, M. Fisher, F. Sadri, and M. Thielscher, Eds., vol. 5405. Springer, 2009, pp. 154–172.

[49] A. Hussain and F. Toni, "Assumption-based argumentation for communicating agents," in *Proceedings of "The Uses of Computational Argumentation", AAAI Fall Symposium*, T. Bench-Capon, S. Parsons, and H. Prakken, Eds. Stanford University, 2009.

[50] J. van Veenen and H. Prakken, "A protocol for arguing about rejections in negotiation," in *Proceedings of the 2nd International Workshop on Argumentation in Multi-Agent Systems (ArgMAS 2005), affiliated to AAMAS 2005*, ser. Lecture Notes in Computer Science, S. Parsons, N. Maudet, P. Moraitis, and I. Rahwan, Eds., vol. 4049. Springer, 2006, pp. 138–153.

[51] H. Prade, "A qualitative bipolar argumentative view of trust," in *Proceedings of the 1st International Conference on Scalable Uncertainty Management (SUM 2007)*, ser. Lecture Notes in Computer Science, H. Prade and V. S. Subrahmanian, Eds., vol. 4772. Springer, 2007, pp. 268–276.

[52] P. Dondio and S. Barrett, "Presumptive selection of trust evidences," in *Proceedings of the 6th International Conference on Autonomous Agent and Multi-Agent Systems (AAMAS'07)*, E. H. Durfee, M. Yokoo, M. N. Huhns, and O. Shehory, Eds., 2007, p. 166.

[53] S. Parsons, P. McBurney, and E. Sklar, "Reasoning about trust using argumentation: A position paper," in *Proceedings of the Seventh International Workshop on Argumentation in Multi-Agent Systems (ArgMAS 2010), affiliated to AAMAS 2010*, 2010.

[54] D. Walton, C. Reed, and F. Macagno, *Argumentation Schemes*. Cambridge University Press, 2008.

[55] J. Sabater and C. Sierra, "Review on computational trust and reputation models," *Artifical Intelligence Review*, vol. 24, no. 1, pp. 33–60, 2005.

[56] C. Castelfranchi and R. Falcone, "Trust is much more than subjective probability: Mental components and sources of trust," in *Proceedings of the 33rd Annual Hawaii International Conference on System Sciences (HICSS-33)*, 2000.

[57] E. Staab and T. Engel, "Combining cognitive with computational trust reasoning," in *Proceedings of the 11th International Workshop on Trust in Agent Societies (AAMAS-TRUST'08)*, ser. Lecture Notes in Computer Science, R. Falcone, K. S. Barber, J. Sabater-Mir, and M. P. Singh, Eds., vol. 5396. Springer, 2008, pp. 99–111.

[58] P.-A. Matt, M. Morge, and F. Toni, "Combining statistics and arguments to compute trust," in *Proceedings of the 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, W. van der Hoek and G. A. Kaminka, Eds., 2010.

[59] B. Yu and M. P. Singh, "Distributed reputation management for electronic commerce," *Computational Intelligence*, vol. 18, no. 4, pp. 535–549, 2002.