Chapter 1
The Added Value of Argumentation

Sanjay Modgil, Francesca Toni, Floris Bex, Ivan Bratko, Carlos I. Chesñevar, Wolfgang Dvořák, Marcelo A. Falappa, Xiuyi Fan, Sarah Alice Gaggl, Alejandro J. García, María P. González, Thomas F. Gordon, João Leite, Martin Možina, Chris Reed, Guillermo R. Simari, Stefan Szeider, Paolo Torroni and Stefan Woltran

---

Sanjay Modgil
Kings' College London, UK, e-mail: sanjay.modgil@kcl.ac.uk

Francesca Toni
Imperial College London, UK, e-mail: ft@imperial.ac.uk

Floris Bex
University of Dundee, UK, e-mail: florisbex@computing.dundee.ac.uk

Ivan Bratko
University of Ljubljana, Slovenia e-mail: bratko@fri.uni-lj.si

Carlos I. Chesñevar
Universidad Nacional del Sur, Bahía Blanca, Argentina, e-mail: cic@cs.uns.edu.ar

Wolfgang Dvořák
Vienna University of Technology, Austria , e-mail: dvorak@dbai.tuwien.ac.at

Marcelo A. Falappa
Universidad Nacional del Sur, Bahía Blanca, Argentina, e-mail: mfalappa@cs.uns.edu.ar

Xiuyi Fan
Imperial College London, UK, e-mail: x.fan11@imperial.ac.uk

Sarah Alice Gaggl
Vienna University of Technology, Austria, e-mail: gaggl@dbai.tuwien.ac.at

Alejandro J. García
Universidad Nacional del Sur, Bahía Blanca, Argentina, e-mail: ajg@cs.uns.edu.ar

María P. González
Universidad Nacional del Sur, Bahía Blanca, Argentina, e-mail: mpg@cs.uns.edu.ar

Thomas F. Gordon
Fraunhofer FOKUS, Berlin, Germany, e-mail: thomas.gordon@fokus.fraunhofer.de

João Leite
New University of Lisbon, Portugal, e-mail: jleite@fct.unl.pt

Martin Možina
University of Ljubljana, Slovenia e-mail: martin.mozina@fri.uni-lj.si

Chris Reed
University of Dundee, UK, e-mail: chris@computing.dundee.ac.uk

Guillermo R. Simari
Universidad Nacional del Sur, Bahía Blanca, Argentina, e-mail: grs@cs.uns.edu.ar

Stefan Szeider
Vienna University of Technology, Austria, e-mail: stefan@szeider.net

Abstract We discuss the value of argumentation in reaching agreements, based on its capability for dealing with conflicts and uncertainty. Logic-based models of argumentation have recently emerged as a key topic within Artificial Intelligence. Key reasons for the success of these models is that they are akin to human models of reasoning and debate, and their generalisation to frameworks for modelling dialogues. They therefore have the potential for bridging between human and machine reasoning in the presence of uncertainty and conflict. We provide an overview of a number of examples that bear witness to this potential, and that illustrate the added value of argumentation. These examples amount to methods and techniques for argumentation to aid machine reasoning (e.g. in the form of machine learning and belief functions) on the one hand and methods and techniques for argumentation to aid human reasoning (e.g. for various forms of decision making and deliberation and for the Web) on the other. We also identify a number of open challenges if this potential is to be realised, and in particular the need for benchmark libraries.

1.1 Introduction

1.1.1 An Overview of Argumentation

The theory of argumentation is a rich, interdisciplinary area of research straddling philosophy, communication studies, linguistics, psychology and artificial intelligence. Traditionally, the focus has been on 'informal' studies of argumentation and its role in natural human reasoning and dialogue. More recently, formal logical accounts of argumentation have come to be increasingly central as a core study within Artificial Intelligence [9], providing a promising paradigm for modelling reasoning in the presence of conflict and uncertainty, and for communication between reasoning entities [1]. In these works, an argument consists of premises and a claim expressed in some logical language $\mathscr{L}$, where the premises support the claim according to some localised notion of proof. For example, the claim that 'Information about Tony should be published' is supported (via application of modus ponens) by the premises: 'Tony has political responsibilities'; 'the information about Tony is in the national interest'; 'if a person has political responsibilities and information about that person is in the national interest then that information should be published'. The

Paolo Torroni
University of Bologna, Italy, e-mail: p.torroni@unibo.it

Stefan Woltran
Vienna University of Technology, Austria, e-mail: woltran@dbai.tuwien.ac.at

[1] As witnessed by the recently inaugurated series of international conferences and workshops www.comma-conf.org, www.mit.edu/~irahwan/argmas/argmas11, www.csd.abdn.ac.uk/~niroren/TAFA-11) and major European research projects (ARGUGRID: www.argugrid.eu, ASPIC: www.cossac.org/projects/aspic, IMPACT: www.policy-impact.eu)

arguments thus constructed are then evaluated in the light of their interactions with other arguments. For example, the preceding argument $A1$ is 'attacked' by the argument $A2$ claiming 'Tony does not have political responsibilities', supported by (because) 'Tony resigned from parliament' and 'if a person resigns from parliament then that person no longer has political responsibilities'. $A1$ therefore loses out at the expense of the winning argument $A2$. Consider the following counter-argument to $A2$: $A3$ = 'Tony does have political responsibilities because Tony is now middle east envoy and if a person is a middle east envoy then that person has political responsibilities'. $A3$ attacks $A2$ by contradicting $A2$'s claim, and $A2$ attacks $A1$ by contradicting a premise in $A1$. The winning arguments can then be evaluated. $A1$ is attacked by $A2$, but since $A2$ is itself attacked by $A3$, and the latter is not attacked, we obtain that $A1$ and $A3$ are the winning arguments.

This example illustrates the modular nature of argumentation that most formal theories (models) of argumentation adopt: 1) arguments are constructed in some underlying logic that manipulates statements about the world; 2) interactions between arguments are defined; 3) given the network of interacting arguments, the winning arguments are evaluated. The appeal of the argumentation paradigm resides in this intuitive modular characterisation that is akin to human modes of reasoning. Also, recent work in AI, and the computer science community at large, has illustrated the potential for tractable implementations of logical models of argumentation, and the wide range of application of these implementations in software systems. Furthermore, the inherently dialectical nature of argumentation models provide principled ways in which to structure exchange of, and reasoning about, justifications/arguments for proposals and or statements between human and/or automated reasoning entities (agents).

Consider the above example where, instead of a single agent engaging in its own internal argumentation to arrive at a conclusion, we now have two agents, Greg and Alistair, involved in a dialogue. Greg proposes $A1$, Alistair $A2$, and then Greg counters with $A3$. This represents a dialogue where each participant has the goal of persuading the other to adopt a belief through the process of exchanging arguments that must interact and be evaluated according to the underlying model of argumentation.

Of course, dialogues introduce an added dimension, in the sense that realistic dialogues often involve more than simply the exchange of arguments. For example, Alistair might challenge a premise in argument $A1$, by asking why the information about Tony is in the national interest. The burden of proof is on Greg to provide an argument as to why this information is in the national interest. Otherwise, Alistair can be legitimately be said to be 'winning' the argument or dialogue. The formal study of dialogue models therefore accounts for a broader range of statements or 'locutions' than simply those involving submission of arguments, as well as the strategic behaviour of interlocutors.

The construction, evaluation and exchange of arguments and related locutions, has great potential for application in the general area of agreement technologies. Arguably, any non-trivial process resulting in an agreement presupposes some kind of conflict and the need to resolve the conflict. Such conflicts may arise between the positions or preferences held by parties involved in negotiating over some kind of resource, or between the beliefs of parties engaged in debate and dialogue, where the purpose is to arrive at some settled (agreed) view. More generally, conflicts will arise whenever alternative outcomes present themselves, independently of whether the parties involved adhere to them or not, for example when parties deliberate over an appropriate course of action from amongst a number of alternatives. In such cases, the alternatives are simply those that present themselves, independently of whether any given party has a particular interest in pursuing a given alternative.

In these dialogues, the reasons or arguments for offers, stated beliefs, or proposed actions can be usefully used to further the goal of the dialogue. The goal of the dialogue may determine a specific set of statements or allowed locutions, as well as rules for making locutions at any point in the dialogue, and rules for determining the outcome of the dialogue. These rules are encoded in a dialogue's protocol. Consider for example the following negotiation dialogue between a buyer and seller of cars in which locutions also involve making, accepting and rejecting offers:

Seller - Offer: Renault
Buyer - Reject: Renault
Buyer - Argue: Because Renault is a French make of car, and French cars are unsafe
Seller - Argue: Renaults are not unsafe as Renaults have been given the award of safest car in Europe by the European Union
Buyer - Accept: Renault

The above example illustrates the utility of argumentation-based models of reasoning and their application to dialogues. Online negotiations involving automated software agents are a key area of research and development. In a handshaking protocol, a seller would simply successively make offers and have these either rejected or accepted. The exchange of arguments provides for agreements that would not be reached in simple handshaking protocols. In the above example, it is by eliciting the reason for the rejection, and successfully countering this reason, that the seller is then able to convince the buyer to buy the car.

The above introduction to argumentation articulates some general reasons for why argumentation may be of value in agreement technologies. In what follows, we more precisely articulate the added value that argumentation brings, above and beyond existing non-monotonic approaches to reasoning in the presence of uncertainty and conflict more generally.

1.1.2 Bridging between Machine and Human Reasoning

Many theoretical and practical developments in argumentation build on Dung's seminal abstract theory of argumentation [29]. A Dung *argumentation framework* (*AF*) consists of a conflict-based binary *attack* relation $\mathscr{C}$ over a set of arguments $\mathscr{A}$. The justified arguments are then evaluated based on subsets of $\mathscr{A}$ that are referred to as *extensions*, and that are defined under a range of semantics. Irrespective of the chosen semantics, the arguments contained in an extension are required to not attack each other (the extensions are *conflict free*), and attack any argument that in turn attacks an argument in the extension (extensions *defend* their contained arguments). Dung's theory has been developed in a number of directions. These include argument game proof theories [76] in which an argument X is shown to belong to an extension under a given semantics, if the player moving X can defend against attacking arguments moved by the player's opponent. Also, several works augment *AF*s with preferences or values [2, 8, 89], attacks on attacks [5, 74], support relations (e.g., [3]), collective attacks (e.g., [16]), those that accommodate numerical information (e.g., [32]), and other extensions.

The continuing development and widespread application of Dung's work can in part be attributed to its level of abstraction. *AF*s are simply directed graphs that can be instantiated by a wide range of logical formalisms; one is free to choose a logical language $\mathscr{L}$ and define what constitutes an argument and attack between arguments defined by a theory. The theory's inferences can then be defined in terms of the claims of the theory's justified arguments, so that the above mentioned argument games can be seen as providing proof theories for the logical formalism. Furthermore, the inference relations of existing logics (with their own proof theories) can be given an argumentation-based characterisation. Thus, as shown in [17, 28, 29, 55], the inferences defined by theories in logic programming and non-monotonic logics (e.g. default, auto-epistemic and defeasible logic), can be defined in terms of the claims of the justified arguments of *AF*s instantiated by arguments and attacks defined by theories in these logics. Dung's theory can therefore be understood as a dialectical semantics for these logics, and the argument games can be viewed as alternative dialectical proof theories for these logics.

The fact that reasoning in existing non-monotonic logics can thus be characterised, testifies to the generality of the principle whereby one argument defends another from attack; a principle that is also both intuitive and familiar in human modes of reasoning, debate and dialogue. Indeed, recent, empirically validated work in cognitive science and psychology supports the latter claim, by proposing that the cognitive capacity for human reasoning evolved primarily in order to assess and counter the claims and arguments of interlocutors in social settings [72].

Argumentation theory thus provides a *language independent* characterisation of both human and logic-based reasoning in the presence of uncertainty and conflict, through the abstract dialectical modelling of the *process* whereby arguments can be moved to attack and defend other arguments. The theory's value can therefore in large part be attributed to its explanatory potential for making non-monotonic reasoning processes inspectable and readily understandable for human users, and it's underpinning of dialogical and more general communicative interactions involving reasoning in the presence of uncertainty and conflict, where such interactions may be between heterogeneous agents (i.e., machine and human). Thus, through such interactions, the reasoning processes of machines can be augmented by intuitive modular argumentation-based characterisations of human reasoning and interaction, and the reasoning processes of humans can be augmented by intuitive modular argumentation-based characterisations of machine reasoning. Indeed, one might argue that the integration of human and machine reasoning is a key requirement for logic-based reasoning techniques to be usefully deployed in practical applications.

It is this value proposition that will be explored in the remainder of this chapter. In Section 1.2 we review some applications and research projects in which human provided arguments, and argumentation-based characterisations of human interactions, are or have been used to inform machine reasoning. In Section 1.3 we review some applications and research projects in which formal models of argumentation are or have or been used to inform human reasoning[2]. Section 1.4 then points towards the need for benchmark libraries for evaluating tools developed for processing Dung frameworks: a key requirement if the value proposition of argumentation is to be realised. Section 1.5 finally concludes.[3]

1.2 Argumentation informing machine reasoning

In this section we review some applications and research projects in which human provided arguments, and argumentative characterisations of human interactions are or have been used to inform machine reasoning. Specifically, these applications incorporate forms of argumentation within:

1. machine learning (in the form of the rule induction CN2 method [23])
2. Dempster-Shafer belief functions [99]

Both approaches use very simple models of argumentation. The first approach uses two types of arguments (attached to examples during the learning phase): positive (to explain/argue why an example is classified as it is) and negative (to explain/argue why an example should not be classified in a certain manner). The second approach uses very simple abstract argumentation frameworks [29] with less than ten argu-

[2] Our reviews in these sections are by no means comprehensive; rather, selected examples are chosen to illustrate the salient points.

[3] Different parts of this chapter have been written/edited by different authors, as follows:

- this Section 1.1 has been written by Sanjay Modgil;
- Section 1.2 has been edited by Francesca Toni, with Section 1.2.1 written by Ivan Bratko and Martin Možina and Section 1.2.2 written by Francesca Toni;
- Section 1.3 has been edited by Sanjay Modgil, with Section 1.3.1.1 written by Sanjay Modgil, Section 1.3.1.2 written by Carlos Chesñevar, Section 1.3.1.3 written by Francesca Toni, QUI Section 1.3.2.1 written by Sanjay Modgil, Section 1.3.2.2 written by Thomas Gordon, Section 1.3.2.3 written by Francesca Toni, Section 1.3.2.4 written by Xiuyi Fan and Francesca Toni, Section 1.3.3 written by Floris Bex, Chris Reed and Sanjay Modgil, and Section 1.3.4 written by Joao Leite and Paolo Torroni;
- Section 1.4 has been written by Wolfgang Dvořák, Sarah Alice Gaggl, Stefan Szeider and Stefan Woltran;
- Section 1.5 has been written by Sanjay Modgil and Francesca Toni.

ments. Despite the simplicity of the underlying argumentation, both approaches give improved performances. We outline these two approaches below.

1.2.1 Argumentation and Machine Learning

1.2.1.1 Overview of Argumentation Based Machine Learning

Machine learning is concerned with the development of algorithms that enable computer programs to learn and improve from experience [73]. The most common type of machine learning (ML) is learning from labeled examples, called also supervised inductive learning. Each example is described by a set of descriptive attributes (inputs), and a class variable (output). The task is to formulate a hypothesis that can infer outputs of examples given inputs. The hypothesis can be used to predict outcomes of new cases, where the true values are unknown.

Machine learning has been shown to be useful in many areas. One of its possible applications is automatic knowledge acquisition to address the bottleneck in building expert systems [41]. While it was shown that it can be successful in building knowledge bases [63], the major problem is that automatically induced models rarely express the knowledge in the way an expert wants. Models that are incomprehensible have less chance to be trusted by experts and other users.

A common view is that a combination of a domain expert and machine learning would be best to address this problem [118]. Most of the applications in the literature combine machine learning and the experts' knowledge in one of the following ways: (a) experts validate induced models after machine learning was applied, (b) experts provide constraints on induced models in the form of background knowledge, and (c) the system enables iterative improvements of the model, where experts and machine learning algorithm improve the model in turns. The last approach is often the most effective; however, it requires considerable effort on the part of the expert. This calls for a method that allows the expert to express his or her knowledge in a most convenient way and combines this knowledge with knowledge extracted from data. In this contribution we discuss argumentation about specific examples as an effective such method. It is commonly accepted that knowledge elicitation based on argumentation, where experts argue about a specific case instead of being asked to articulate general knowledge, is considerably simpler due to the following:

- When providing their knowledge, domain experts have to focus on a specific problem only and do not need to be concerned whether their provided knowledge given for this problem is generally accepted for all possible problems. Counter-arguments will take care of exceptions.
- Disagreements between domain experts do not pose a problem; all provided arguments (for and against) can be imported in the knowledge base and it is left to the reasoner to select which of them are acceptable.

The idea of argument-based machine learning (ABML) [83], a combination of machine learning and argumentation, is to induce a hypothesis that is consistent with learning data and provided arguments. The motivation for using arguments in machine learning lies in two expected advantages:

1. Arguments impose constraints over the space of possible hypotheses, thus reducing overfitting and guiding learning algorithms to induce better hypotheses.
2. An induced theory should make more sense to experts as it has to be consistent with the given arguments provided by the experts.

With respect to advantage 1, by using arguments, the computational complexity associated with search in the hypothesis space can be reduced considerably, and enable faster and more efficient induction of theories. The second advantage is crucial for building knowledge bases. From the perspective of a machine learning method, there are several possible hypotheses that explain the given examples sufficiently well with respect to predictive accuracy; however, some of those hypotheses can be incomprehensible to experts. Using arguments should lead to hypotheses that explain given examples in similar terms to those used by the expert.

During the process of interactive knowledge acquisition with experts and machine learning, it is not rare that provided arguments contradict the data. In such cases, the experts need to either: a) revise their knowledge about the domain, or b) make amendments to the data. Whatever option they decide to choose, both are useful for them. In the first case, they learn something new about the domain, while in the latter, the corrections result in more accurate data.

In ABML, arguments are provided by human experts, where each argument is attached to a single learning example only, while one example can have several arguments. There are two types of arguments; positive arguments are used to explain (or argue) why a certain learning example is in the class as given, and negative arguments are used to explain why it should not be in the class as given. Examples with attached arguments are called *argumented examples*.[4]

An ABML method is required to induce a theory that uses given arguments to explain the examples. If an ABML method is used on normal examples only (without arguments), then it should act the same as a normal machine learning method. We developed the ABCN2 [83] method, which was used in all case-studies described in the following section. ABCN2 is an argument-based extension of the well known method CN2 [23], that learns a set of unordered probabilistic rules from argumented examples. In ABCN2, the theory (a set of rules) is said to explain the examples using given arguments, when there exists at least one rule for each argumented example that is consistent with at least one positive argument (contains argumentative in its condition part) and is not consistent with any negative argument.

[4] Due to space limitations, we will only roughly describe ABML (see [83] for precise details).

1.2.1.2 Interaction between an expert and ABML.

It is not feasible for an expert to provide arguments for all the examples. Therefore, we use the following loop to pick out the *critical examples* that should be explained by the expert. The loop resembles an argument-based dialogue between a computer and an expert.

1. Learn a hypothesis with ABML using given data.
2. Find the most critical example and present it to the expert. If a critical example can not be found, stop the procedure.
3. The expert explains the example; the explanation is encoded in arguments and attached to the learning example.
4. Return to step 1.

A critical example (step 2) is an example the current hypothesis can not explain very well. The hypothesis assigned a wrong class value to this example, and therefore asks the expert to argue why he or she believes this example should be in a different class. Using expert's arguments, ABML will sometimes be able to explain the critical example, while sometimes this will still not be entirely possible. In such cases, we need additional information. The whole procedure for one-step knowledge acquisition (step 3) is described with the next 5 steps:

Step 1: Explaining critical example. The experts are asked the following question: ”Why is this example in the class as given?” Then, the experts provide a set of arguments $A_1,\dots,A_k$ all confirming the example's class value.

Step 2: Adding arguments to example. Arguments A_i are given in natural language and need to be translated into domain description language (attributes). Each argument supports its claim with a number of reasons. When reasons are some attribute values of the example, then the argument can be directly added to the example. On the other hand, if reasons mention other concepts, not currently present in the domain, these concepts need to be included in the domain.

Step 3: Discovering counter examples. Counter examples are used to spot whether the availabe arguments suffice to successfully explain the critical example or not. If ABML fails to explain the example, then the counter examples will show where the problem is. A counter example has the opposite class of the critical example, however arguments given for the critical example apply also for the counter example.

Step 4: Improving arguments. The expert needs to revise the initial arguments with respect to the counter example. This step is similar to steps 1 and 2 with one essential difference; the expert is now asked ”Why is critical example in one class and why counter example in the other?” The answer is added to the initial argument.

Step 5: Return to step 3 if counter example found.

1.2.1.3 Examples applications of ABML

We above outlined ABML, a generic method for integrating argumentation and machine learning. We now give some example scenarios where ABML has been applied.

Construction of Sophisticated Chess Concepts.

For the purposes of a chess tutoring application developed by Sadikov et al. [97], we used ABML to acquire knowledge for two sophisticated chess concepts: *bad bishop* and *attack on king*. In this section, we will shortly discuss the process of knowledge acquisition in both cases and give an overview of the results ([81] and [82] give a more elaborate description of case-studies).

In the bad bishop case, 200 chess positions were selected. For each of them, the experts gave a qualitative assessment whether the bishop in the position was strategically bad or not. We furthermore described the positions with 100 positional features, which served as attributes. These features are commonly used by strong chess programs and suffice for playing chess on a strong level. We used the ABCN2 method to induce a set of rules with the following structure:

IF conjunction of some features THEN *bishop=bad*.

The ABML based knowledge acquisition process discovered eight critical examples and experts explained them with arguments. During the argumentation, experts used five concepts that were not included among 100 default attributes. These five concepts were encoded as five new attributes. Surprisingly, the final rules considered only these five attributes and dropped others that are otherwise very useful for computer play. This demonstrates, on the one hand, how chess players think differently from computers and, on the other hand, suggests that without knowledge introduced through arguments, learned rules would be incomprehensible to experts.

The final model, after all iterations, was evaluated on the test dataset. The improvement of the model was evident: from the initial 72% classification accuracy (Brier score 0.39, AUC 0.80), the final 95% accuracy (Brier score 0.11, AUC 0.97) was achieved.

Our domain experts (a chess master and a woman grandmaster) clearly preferred the ABML approach to manual knowledge acquisition. They tried to formalize the concept of bad bishop without ABML, however it turned out to be beyond their practical ability. They described the process as time consuming and hard, mainly because it is difficult to consider all relevant elements. However, with ABML and by considering only critical examples, the time of experts' involvement decreased, making the whole process much less time consuming.

In the second experiment, involving conceptualization of *attack on king* concept, the process took much longer: 38 iterations. This probably happened because the concept itself is considerably more complicated. The process itself was similar to the one with bishops, with one important difference, the expert changed the class value of positions in 10 out of 38 iterations. In all of those cases they decided to change the class value, as they were unable to argue why they assigned the original class in the first place.

After the ABML process, special care was given to examine the interpretability of rules. The experts compared rules obtained with and without arguments. In the case without arguments, they identified three rules (out of 12) that contained counter-intuitive terms for a chess expert. It is not uncommon for ML to produce such seemingly nonsensical explanations as an artefact of the data. On the other hand, ABML produced 16 rules, and none of them included any illogical terms as deemed by the experts. As our goal is to use this model in a chess tutoring application, such terms could be very harmful. A teacher using illogical argumentation (even of a correct decision) is never a good idea. And it is surprising how harmful are the three rules with illogical terms in our case. With a statistical experiment, we showed that in 85% of the cases, where the model correctly predicted the class value, it used the wrong argumentation to explain its decision.

Acquisition of Neurological Knowledge.

In the following, we will briefly describe the process of knowledge acquisition for a neurological decision support system [56]. Our goal was to learn a rule-based model that would help the neurologists differentiate between three types of tremors: Parkinsonian, essential, and mixed tremor (co-morbidity). The system is intended to act as a second opinion for the neurologists. Our data set consisted of 67 patients diagnosed and treated at the Department of Neurology, University Medical Centre Ljubljana.

Due to a small number of cases, we shall focus only on a qualitative evaluation. Although the final model (after argumentation) had a better accuracy, the small number of available cases limits us from drawing any statistically significant conclusions. For a qualitative evaluation, the domain expert was asked to evaluate each rule according to its consistency with his domain knowledge. We found a significant difference between the evaluation of initial and final rules. All the rules in the final model were consistent with domain knowledge, while three of the starting rules were not. Furthermore, five of the final rules were marked as *strong rules* meaning that they are sufficient for making a diagnosis. In the initial set, the machine learning algorithm identified only one such rule. Moreover, the relevance of the argumentation process involving ABML and expert (with critical and counter examples) was also reflected by the fact that they assisted the expert to spot 2 mistakes in the initial diagnosis. Therefore, such a tool could be a useful addition to their usual practice.

1.2.1.4 Discussion: Open Issues and Challenges

The above experiments demonstrate the benefits that argumentation brings to machine learning. From the perspective of argumentation, there are two sets of open questions that could further improve the synergy between machine learning and argumentation. The first set concerns the type of arguments applicable in ABML.

At the moment, we consider only arguments that directly argue about the outcome of the example: positive arguments and negative arguments rebut each other. The question is, could we also use arguments that undermine (rebut on the *premises* of) other arguments ? An extension of the basic ABML theory considering other types of arguments is given in [80], however, it still needs to be evaluated on practical examples. Furthermore, would it be possible to exploit the structure of argument-based reasoning in ABML? In other words, is it possible to use an attack graph (i.e., Dung framework) of arguments instead of just single arguments? These are some of the ideas that could further increase the added value of argumentation in machine learning.

The second set of questions is related to how the output of ABML methods could help argumentation. Could an ABML method be used to facilitate the construction of a knowledge base for an argumentation-based expert system? Such a method would try to discover rules that would together with an argumentation reasoning mechanism (e.g [29]) infer correct classes for all learning examples. It is unlikely that we are able to learn such rules with ABCN2, as ABCN2 is specialized in learning classification rules. A possible direction would be to interface the ABILP algorithm [18], an argument-based version of induction logic programming (ILP), with argumentation reasoning. Such an ILP algorithm would, instead of classical monotonic reasoning, use non-monotonic reasoning to evaluate candidate hypotheses.

1.2.2 Argumentation and Dempster-Shafer belief functions

1.2.2.1 Overview of Integration

Dempster-Shafer belief functions [99] provide a generalization of the Bayesian theory of subjective probability based on two core concepts: that degrees of belief for one question can be obtained from subjective probabilities for a related question, and a rule for combining these degrees of belief when they are based on independent items of evidence. Yu and Singh [121] deploy Dempster-Shafer belief functions to answer the question of whether a given agent (the evaluator) should trust another (the target), given statistical information concerning the past behaviour of the target. Matt et al [67] integrate argumentation into this method, by proposing a method 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. For concretely evaluating these method, the arguments are built from current and past contracts between evaluator and target (see Section 1.2.2.2 below). Here, we briefly review how argumentation can contribute to defining Dempster-Shafer belief functions to reason about trust.

In general, a belief functions $Bel : 2^{\Omega} \rightarrow [0,1]$, where Ω is a given universe, need to be defined via some evidence mass function, $m : 2^{\Omega} \rightarrow [0,1]$, which needs to be positive, normalised and such that $m(\emptyset) = 0$. Given such m, for every subset

$E \subseteq \Omega$, $Bel(E) = \sum_{X \subseteq E} m(X)$. Yu and Singh [121] use a (Dempster-Shafer) belief function as a mathematical model of trust, where $\Omega = \{T, \neg T\}$ is a simple universe with T ($\neg T$) representing that the evaluator considers the target to be trustworthy (untrustworthy, respectively). In their approach, the evidence mass function may be derived either from the knowledge of the evaluator's own past interactions with the target (local trust rating), or by combination of belief functions representing testimonies from other entities concerning the target (belief combination). Matt et al [67] focus only on local trust rating. In this case, the evidence mass function is defined in terms of the history of past interactions between the evaluator and the target (assuming that this is sufficiently long), classified as *poor*, *satisfying*, or *inappreciable*. Given that the total number of past interactions is $N = N^- + N^+ + N_?$ with N^-, N^+, $N_?$ the number of times the quality of the interaction was poor, satisfying and inappreciable, respectively, the evidence mass function m is given by

$$m(\emptyset) = 0 \quad m(\{T\}) = \frac{N^+}{N} \quad m(\{\neg T\}) = \frac{N^-}{N} \quad m(\Omega) = \frac{N_?}{N}$$

The evaluator can use the belief function obtained from this evidence mass function to decide whether to interact with the target if and only if its trust in the target (i.e. $Bel(\{T\})$) exceeds its distrust (i.e. $Bel(\{\neg T\})$) by a threshold value $\rho \in [0,1]$ that represents how *cautious* the evaluator is.

Matt et al define a new (Dempster-Shafer) belief function Bel_a taking into account, in addition to the statistical information (N^-, N^+, $N_?$), also an abstract argumentation framework F including, amongst its arguments, a set A of arguments each supporting one of T (in favour of trusting the target) or $\neg T$ (against trusting the target). They use F and A to define $\hat{p}_A : 2^\Omega \rightarrow [0,1]$, the *argumentation-based prior* as

$$\hat{p}_A(E) = \frac{1}{I}\left[\hat{p}(E) + V_A \sum_{a \in A} s_F(a)\hat{p}(E|\{X_a\})\right]$$

where (see [67] for details):

- $\hat{p}(E)$ is the statistical prior, determined from N^- and N^+
- I and V_A are suitably defined parameters, informally representing the total amount of information available (I) and the information contributed by arguments in A (V_A), namely how much arguments for or against trust count in determining trustworthiness, in relation to statistical information
- $\hat{p}(E|\{X_a\})$ is the conditional probability of E given the conclusion X_a of argument $a \in A$
- $s_F(a)$ gives the strength of argument $a \in A$; this strength is measured taking into account all arguments in F (and not solely those in A within F) as well as the attack relation amongst arguments

Then, Bel_a is obtained, according to the standard Dempster-Shafer theory, from the *argumentation-based evidence mass function* $m_A : 2^\Omega \rightarrow [0,1]$ given by

$$m_A(\emptyset) = 0 \qquad m_A(\{T\}) = (1-\varepsilon_A)\,\hat{p}_A(\{T\})$$
$$m_A(\Omega) = \varepsilon_A \qquad m_A(\{\neg T\}) = (1-\varepsilon_A)\,\hat{p}_A(\{\neg T\})$$

with ε_A a parameter giving a measure of the uncertainty of the evaluator given the past interactions with the target (see [67] for details).

Although defined in the context of trust computing, this method for combining argumentation and statistics is generic, in that $\hat{p}_A$ (and thus Bel_a) can be obtained for any given argumentation framework F with special arguments A (for answering a question) given a prior $\hat{p}$.

From an argumentation perspective, this method requires a way to compute the (numerical) strength of arguments in an abstract argumentation framework. This could be defined as 1 for "acceptable" arguments according to some argumentation semantics (e.g. admissibility as in [29]) and 0 otherwise, or according to some quantitative notion, e.g. presented in one of [13, 20, 68]. In the experimental evaluation of this model for trust, discussed later on in Section 1.2.2.2, the quantitative, game-theoretic notion of strength of [68] is considered. Furthermore, in Section 1.2.2.3, we outline some open issues for deploying argumentation-based belief functions for trust and in general.

1.2.2.2 Arguments from contracts for trust computing

The above described method integrating abstract argumentation and (Demptster-Shafer) belief functions has been applied in the context of assessing trust in contract-regulated interactions in general [67], but with emphasis on interactions amongst service providers and service requestors in service-oriented architectures, with contracts represented by SLAs (Service Level Agreements). In this setting, the argumentation framework F consists of arguments for or against trust (A), based upon the existence or lack (respectively) of contract clauses providing evidence for one of four dimensions or service provision (namely availability, security, privacy and reliability). In addition, F may also contain up to four arguments (in $F \setminus A$) attacking an argument for trust along a dimension on the ground that the target has in the past "most often" violated existing contract clauses concerning that dimension.

Matt et al evaluate their method experimentally, in this service-oriented setting, against the method of Yu and Singh [121], relying upon statistical information only (see Section 1.2.2). The two methods have identical predictive performance when the evaluator is highly "cautious", but the use of arguments built from contracts gives a significant increase when the evaluator is not or is only moderately "cautious". Moreover, target agents are more motivated to honour contracts when evaluated using the argumentation-based model of trust than when trust is computed on a purely statistical basis.

1.2.2.3 Summary and open issues

In conclusion, the method integrating abstract argumentation and (Demptster-Shafer) belief functions has been applied in the context of assessing trust in contract-regulated interactions, and in particular in the setting of service-oriented architectures. However there a number of open challenges. Firstly, the experimental setting makes use of a limited set of arguments (based upon the existence or lack of contracts and the tendency of target agents to default their contractual commitments); it would be interesting to consider a broader set of arguments, e.g. taking into account opinions by other agents. Secondly, the method has been experimented with in a simulated environment, whereas it would be interesting to apply it in a real setting. Finally, although defined in the context of trust computing, the method for combining argumentation and statistics is generic, as discussed earlier; it would be interesting to study further applications of this method, to see how useful and effective it is.

1.3 Argumentation informing human reasoning

In this section we review a number of applications and research projects in which formal models of argumentation are or have been used to inform human understanding, reasoning and debate . Specifically, these applications utilise one or more of the following:

1. Models structuring the contents of individual arguments, and the way in which these contents are related, have been used in explaining the reasoning of machines.
2. Models of the dialectical relationships between arguments have been used to guide authoring and mapping of arguments by individuals and in debates and opinion gathering forums, and evaluate the status of arguments.
3. Formal dialogical models have been used to mediate the rational exchange of arguments between humans and/or automated agents.

In Section 1.3.1 we briefly review the use of models of argument for structuring explanations in medical decision making tools, and then go on to discuss more recent uses of argumentation in decision making. Section 1.3.2 then considers the use of argumentation in distributed decision making, in which participants exchange arguments for and against proposals for action. Specifically, we review previous and current European Union funded research on development of tools for facilitating distributed decision making. Some of these make use of the schemes and critical questions approach to structuring arguments and their interactions [115] that is key to facilitating the use of argumentation in guiding rational and focussed deliberation. We also review, in Section 1.3.2.4, some recent work on argumentation-based dialogues, that can be used to support deliberation as well as several other forms of exchanges in distributed settings. Finally, the plethora of existing argument visualisation and mapping tools [62] (e.g., Rationale [10] and Araucaria [94]) testifies to the enabling function of argumentation models in guiding rational human reasoning and debate. A number of these tools are available online suggesting the notion of an *argument web* in which authored arguments can be exchanged and reused. It is in the context of this envisaged *argument web* that Section 1.3.3 reviews recent work on tools for argument mapping and authoring. Section 1.3.4 then suggests how argumentation can enhance interaction in the *social web*.

1.3.1 Argumentation-based Decision Making

1.3.1.1 Medical Decision Making.

Amongst the earliest works that utilise formal model of argument, are the medical expert systems developed by researchers at Cancer Research UK (see [43] and `www.cossac.org/projects/archive`). A key feature of these applications is that knowledge resources, augmented by human entered data, are used in making some recommendation. The reasoning by which these recommendations are made, are presented in the form of arguments for and against the recommendations. For example, the *REACT* system [48] supports a doctor's consultation with a patient at risk from ovarian or breast cancer. The system visually shows how risk levels are affected by combinations of various medical interventions and other planned patient decisions (e.g., having a baby), where these changes in risk are evaluated using rules encoded in the system. A key explanatory function of the system is the presentation of arguments for and against a given intervention, where these arguments (justifying a reduction/increase in risk) are constructed based on the aforementioned rules, and are augmented by other arguments relevant to the well being of the patient. A key feature is that the structuring of individual arguments is based on the Toulmin model of argument structure [111], whereby an argument consists of a claim (e.g., remove ovaries) justified by given data (the patient is over 40 and a BRCA2 gene carrier) and a warrant linking the data to the claim (patients over 40 who are BRCA2 gene carrier are reasons to remove ovaries for prevention of cancer), supported by a *backing* (the clinical studies that support the warrant) and with some qualifier indicating the strength of the claim (the degree of risk reduction).

1.3.1.2 Dialectical Explanation for Decision Making.

Recent work by Argentinean researchers in argumentation has led to formalizing and implementing several aspects of argumentation for decision making. In [46] the concept of *dialectical explanation* was introduced and can be applied for decision making domains. The purpose of a dialectical explanation is to transfer the understanding of how the warrant status of a particular argument can be obtained from a given argumentation framework. When applying this framework in a decision making domain, the dialectical explanation can provide, as formulated in [90], an advice that can be presented in a form which can be readily understood by the decision maker; and since that explanation reflects the argumentative analysis that was carried out, it provides access to both the information and the reasoning that underpins the given advice. In [42] a model for *defeasible* decision making was introducing by combining *defeasible* decision rules and arguments. The principles stated in that work were exemplified in a robotic domain, where a robot should make decisions about which box must be transported next. In that decision framework, the agent's decision policy can be changed in a flexible way, with minor changes in the criteria that influence the agent's preferences and the comparison of arguments. The proposal includes a simple methodology for developing the decision components of the agent.

Providing a full-fledged model for characterizing explanation in decision making involves a number of open issues and challenges. Significant research has been dedicated to the enhancement of the explanation capabilities of knowledge-based systems and decision support systems, particularly in *user support systems*. Recent investigations have shown how to enhance them using argumentation techniques [21] for providing rational recommendations supported by a procedure explicitly justified. An open issue is the integration of quantitative and qualitative information when providing explanations, so that the systems can perceived as more reliable and user-friendly. The strength of an explanation can also be affected by the existence of several arguments supporting a given conclusion (*i.e.*, argument accrual). New argument-based inference procedures for the accrual of arguments have been developed [65], but their deployment in actual *Argument-based Decision Support Systems* (ArgDSS) requires further investigation.

Another interesting aspect for decision making concerns the development of so called *Argument Assistance Systems* (AAS) [113] and *Hybrid Argument Systems* (HAS) [52]. While AAS focus on graphical-oriented functionalities for graphically representing an argumentation process (providing facilities for creating and analyzing arguments and their interrelationships), HAS aim to combine such facilities with an automatic inference procedure for evaluating the argumentation semantics under consideration. Following these ideas, some ArgDSS implementations have explicitly considered usability [50]. However, there are no standard adopted model and criteria for assessing the usability of ArgDSS within the argumentation community, mainly due to the necessity of developing interfaces of a novel kind in an area where there is still much to be learnt about the way arguments can be sensibly and clearly presented to the users [113]. It is necessary to further explore alternative usability-oriented evaluations to validate and improve the usability-oriented design guidelines currently identified, as well as the corresponding usability principles in play. In particular, the datamining technique presented in [51] for detecting and charaterizing common usability problems of particular contexts of usage (such as ArgDSS) is under consideration. For the particular case of the DeLP (Defeasible Logic Programming) Client that interacts with a DeLP Server [44] which provides a reasoning service, an incremental iterative usability-oriented development process is being performed. In the near future, direct manipulation of arguments has to be considered, leading to a novel interaction style for ArgDSS as well as the revision of the questions associated with every design guideline to cover it.

A key challenge for development of argumentation based decision support systems, concerns a key concept mentioned by Girle et al. [90], which refers to the detailed analysis of the epistemic state of the decision maker, providing a suitable model for considering or obtaining those salient features (”unusual details” in their words) that might led to alternative decisions. Such details can be introduced as triggers of changes in the beliefs to adapt the agent's epistemic state when considering the acceptance of a new piece of information as part of that state.

Argumentation research impacts belief revision research by introducing consideration of the support of each belief in the epistemic state; this support has the form of arguments that can be constructed from that state. Each belief takes the role of the claim of an argument built from a set of premises, and the decision of accepting the belief is made after considering the status of all the arguments in favor of and against the argument supporting the claim.

Investigating the multifaceted relationship between Belief Revision and Argumentation requires considering cross-links between different aspects on either side while also considering their place in the higher context of reasoning. There has been recent work trying to define change operations on argumentation frameworks [38, 39]. Among them, we may group those defining revision operators and those defining contraction operators. For instance, in [96, 78, 79] revision operators are defined in order to warrant some (new) claim, and in [47] different contraction operators are defined in order to retract some inferences from the original knowledge base.

Further steps exploring the relation between argumentation and the dynamics of beliefs are necessary. An interesting area to explore is the one dedicated to decision support systems dedicated to diagnosis. For instance, if I is a query such as “If element α is supplied, will effect β be produced?”, whatever the element α and the effect β are, reasoning will become hypothetical to answer the *what-if* query; this type of query will require the consideration of alternative hypothetical epistemic states, a complex task whose outcome could be improved combining belief revision and argumentation.

1.3.1.3 Decision Making in ARGUGRID.

The ARGUGRID project (funded by the EC, 2006–2009) [5] developed a platform populated by rational decision-making agents associated with service requestors, service providers and users [106], to be used in the context of grid and service-oriented applications. Within agents, argumentation as envisaged in the Assumption-Based Argumentation (ABA) framework [17, 25, 27, 26] is used to

[5] `www.argugrid.eu`

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 [69, 70, 31]. Here, argumentation is used to compute "optimal" decisions, in ways that have a direct correspondence in standard, *normative* decision theory. For example, the method in [70] computes dominant decisions, and the method in [31] deploys the minimax principle. The use of argumentation, however, also provides a *descriptive* explanatory counter-part to the optimal decisions. An overview of the decision-making methods deployed in ARGUGRID can be found in [105].

The ARGUGRID approach to decision making has been validated by way of industrial application scenarios in e-procurement and earth observation [106, 69, 70] (as described later in Section 1.3.2.3).

1.3.2 Argumentation-based Agreement

In Section 1.1 we described how theories of argumentation have provided a basis for development of dialogical models supporting the exchange of information in order to arrive at an agreement. In particular, there have been proposals for generalising argumentation-based decision making to cases where multiple (human and or automated) agents deliberate to agree on a preferred course of action. To illustrate, we briefly review the *CARREL* system [101, 103, 102], developed as part of the European Union *ASPIC* project on argumentation models and technologies[6] to support the exchange of arguments across several agents, the approach to deliberative democracy taken in the current IMPACT project[7] and the approach to inter-agent negotiation developed within the ARGUGRID project[8]. Finally, we overview a generic form of argumentation-based dialogue to support agreement by means of various forms of dialogues, ranging from information-seeking to deliberation.

1.3.2.1 CARREL.

The CARREL system developed a dialogue manager that mediated the exchange of arguments between geographically distributed human agents deliberating over whether a given available organ was viable for transplantation to a given recipient. The aim of the system was to increase the likelihood that an organ would be transplanted, in cases where the medical guidelines suggested the organ was unsuitable, but a well argued case for deviating from the guidelines could be made. One of the main challenges in developing the system was to realise the key aim of using argumentation-based models to facilitate rational reasoning and debate. To this

[6] `www.cossac.org/projects/aspic`
[7] `www.policy-impact.eu`
[8] `www.argugrid.eu`

end, CARREL made extensive use of the schemes and critical questions (*ScCQ*) approach [115, 117]. For the moment we digress from our description of of CARREL to explain how *ScCQ* can be used to bridge between formal models and human argumentation.

Argument schemes identify generic patterns of reasoning that can be represented as rules in formal logic or as natural language templates. These generic argument schemes (upwards of sixty have been identified) are then associated with critical questions that identify the presumptions that any specific instantiation of the scheme (i.e., an argument) makes, and thus the potential points of attack by counter-arguments that may themselves be instances of argument schemes with their own critical questions. For example, consider [4]'s argument scheme – *SA* – for action:

> 'In circumstances S, action A achieves goal G which promotes value V, and so action A should be done'.

The variables in this scheme can be instantiated by a human or logic based agent $Ag1$ (in which case the scheme would be represented as a defeasible implication) to construct a specific argument $Arg1$, where S = 'Saddam has weapons of mass destruction (wmd)', A = 'invade Iraq', G = 'remove wmd' and V = 'world peace'. Critical questions for the scheme *SA* include, 'Is S the case ?', 'Does G promote V ?', 'Are there alternative actions for realising G', e.t.c. Each of these questions can then be addressed by an agent $Ag2$ as a question in its own right, so placing the burden of proof on $Ag1$ to justify the questioned presumption with a supporting argument that might itself be an instance of a scheme. A question can also be addressed as a counter-argument instantiating a scheme. For example, consider the scheme – *SE* – from expert opinion:

> 'E is an expert in domain D, E asserts that A is known to be true, A is in domain D, and so A is true'.

$Ag2$ might instantiate this scheme with E = 'Hans Blick', A = 'Saddam does not have wmd', D = 'weapons inspection', yielding an argument $Arg2$, which instead of *questioning* the premise 'Saddam has wmd' in $Arg1$, *attacks* $Arg1$ on this premise. *SE* has its own critical questions (e.g., 'is E an expert in domain D ?), and so $Arg2$ can be attacked by arguments (instantiating schemes) addressing these critical questions, and so on. In general then, one can see that schemes and critical questions can be used to guide rational exploration through a space of possible argumentation, providing for human and machine authoring of arguments, and identification of relevant counter-arguments.

In employing the *ScCQ* approach, the developers of CARREL realised the need for schemes and critical questions that were more tailored to the domain of organ transplantation, in order to effectively guide argument-based deliberation over the viability of organs. The development of this tailored set of *ScCQ* was undertaken in consultation with domain experts. The implemented CARREL dialogue manager was then deployed to animate these specialised *ScCQ*, presenting arguments to agents, together with their associated critical questions, and the schemes that could be used to address these questions as attacking arguments. The arguments exchanged during the course of a deliberation, were then organised into a Dung argumentation framework, and together with sources of knowledge providing information about the relative strengths of (preferences over) arguments, the frameworks were evaluated to determine whether an argument assigning an organ to a recipient was winning (see Section 1.1.2).

1.3.2.2 IMPACT.

CARREL was intended primarily for use by human (medical) experts. On the other hand, the current *IMPACT* project intends to engage both experts and lay members in policy deliberation. IMPACT is a three year European Union project[9] that began in 2010, and aims to develop and integrate formal, computational models of policy and arguments about policy, to facilitate deliberations about policy at a conceptual, language-independent level.

The basic idea of deliberative democracy is to empower citizens with the means to participate in a more direct way in the development and evaluation of policy alternatives. However, the current state-of-the-art in eParticipation technology, in which arguments are exchanged in natural language using discussion forums, weblogs and other social software, cannot scale up to handle large-scale policy deliberations, as it requires too much manual translation, moderation and mediation to be practical. As the number of participants increases, it becomes more and more difficult for participants to follow the discussion and keep track of the issues and arguments which have been made, even when they are fluent in the language, not to mention messages in foreign languages. The signal-to-noise level in discussion forums can be very low, due to repetition of points which have already been made, personal attacks and other ad hominem arguments, by persons who are more interested in provoking others or attracting attention to themselves than in constructively contributing to a rational debate.

The IMPACT project thus aims to apply state-of-the-art argumentation technology to facilitate more rational, focussed, deliberative forms of democracy. Specifically, the phases of a policy cycle can be sequenced as: 1) agenda setting, 2) policy analysis, 3) lawmaking, 4) administration and implementation, and 5) monitoring. IMPACT is focusing on the second policy analysis phase. The project aims to:

1. Develop argument schemes and critical questions specifically orientated towards deliberation and debate about policy, and to use these *ScCQ* to automatically generate online surveys that invite lay members of the public to submit their opinions. The guidance provided by the *ScCQ* will overcome many of the problematic issues highlighted above.
2. Leverage the explanatory capabilities of argumentation-based structuring of knowledge. IMPACT is using methods from the field of Artificial Intelligence and Law to model policies as context-dependent rules or principals which may conflict with one another or be subject to exceptions and to simulate the effects of these

[9] `www.policy-impact.eu`

policies on a range of cases, using an inference engine based on a computational model of argumentation. The policy models built with these tools will improve the ability of citizens and government to predict the impact of policy measures on both specific cases and on an aggregated set of benchmark cases as a whole. For example, models of social benefits or tax policy of this kind would enable citizens to predict the impact of proposed policy changes on their entitlements or tax burden, respectively.

3. Provide tools for experts to mine arguments from natural language text so enabling the huge amounts of information publicly available on the Internet (for example in web sites, online newspapers, blogs and discussion forums) to be intelligently harvested to gather stakeholders' interests, values, issues, positions and arguments about policy issues. More specifically, such tools would enable the vast amount of information available on public sector resources on the Internet to be optimally used and reused in policy deliberations.
4. Develop dialogical models and software methods and tools for constructing, evaluating, and visualizing arguments to meet the challenges of large-scale public deliberations on the Internet.

1.3.2.3 Negotiation in ARGUGRID.

In the ARGUGRID platform [106], argumentation is used, in a grid/service-oriented architecture setting, to support the negotiation between agents [31, 60] on behalf of service requestors/providers/users, as well as to support decision making (as described in Section 1.3.1.3). This negotiation takes place within dynamically formed virtual organisations [71]. The agreed combination of services, amongst the argumentative agents, can be seen as a complex service within a service-oriented architecture [104].

The need for negotiation arises when agents have conflicting goals/desires but need or may benefit from cooperation in order to achieve them. In particular, this cooperation may amount to a change of goals (e.g. towards less preferable, but socially acceptable goals) 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).

Argumentation-based negotiation is a particular class of negotiation, whereby agents can provide arguments and justifications as part of the negotiation process [61]. It is widely believed that the use of argumentation during negotiation increases the likelihood and/or speed of agreements being reached [91]. In ARGUGRID, argumentation, in the form of ABA [17, 25, 27, 26], was used to support negotiation between a buyer and a seller (e.g. of services) and resulting in (specific forms of) contracts, taking into account contractual properties and preferences that buyer and seller have [31]. Here, negotiation is seen as a two-step process, with a first step where ABA is used to support decision making (see section 1.3.1.3), and then a second step uses a *minimal concession strategy* [31] 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 for the sake of reaching agreement. This strategy can also incorporate rewards during negotiation [30], where rewards can be seen as arguments in favour of agreement.

ABA is also used to support negotiation in [60], for improved effectiveness, in particular concerning the number of dialogues and dialogue moves that need to be performed during negotiation without affecting the quality of solutions reached, in more general resource reallocation settings. This work complements studies on protocols for argumentation-based negotiation (e.g. [112]) and argumentation-based decision making during negotiation as discussed earlier for [31], by integrating argumentation-based decision making with the exchange of arguments to benefit the outcome of negotiation. In this work, 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. The benefits of providing these reasons are assessed 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.

We conclude by listing three main scenarios in which ARGUGRID applied argumentation-based methods for decision making, negotiation and trust computing (see [106] and `www.argugrid.eu` for details):

- e-procurement [69], in particular for an e-ordering system, where service providers sell e-procurement products and service requestors are users needed a combination of these products to fulfil their goals;
- earth observation [70], in particular for checking oil-spills, where service providers return or manipulate images (e.g. from satellites) and service-requestors are users need (processed) images to fulfil their goals;
- e-business migration [30], investigating the development of formal frameworks for modelling contracts, contract negotiation and conflict resolution that are essential in the business process for outsourcing activities, focusing on a migration of computer assembly activities setting.

1.3.2.4 Argumentation-based Dialogues

Argumentation-based dialogue systems have attracted substantial research interest in the recent years. In [88], Prakken has presented a brief summary of the development of dialogue systems. The modern study of formal dialogue systems for argumentation starts from Charles Hamblin's work [57]. The topic was initially studied within philosophical logic and argumentation theory [66, 116]. Subsequently, researchers from the field of artificial intelligence & law [53, 87] and multi-agent systems [1, 120] have looked into dialogues systems as well.

Two major questions need to be addressed in a study of dialogue models. Firstly, how to construct "coherent" dialogues? Secondly, how to construct dialogues with specific goals? The first question is addressed by introducing *dialogue protocols*; and the second question is addressed by studying *dialogue strategies*.

A more recent effort in formalising two-agent dialogues can be seen in [40]. In this work, Fan and Toni define a dialogue protocol for generic dialogues. They have used Assumption-based Argumentation (ABA) [26] as the underlying representation, as ABA is a general-purpose, widely used argumentation framework. In their model, a dialogue is composed of a sequence of utterances of the form

$$\langle From, To, Target, Content, ID \rangle,$$

where *From* and *To* are agents; *Target* and *ID* are identifiers; and *Content* is either a topic, a rule, an assumption, a contrary[10], or pass. A dialogue starts with an agent posing a topic and completes when both agents utter pass.

To ensure the integrity of a dialogue, Fan and Toni have introduced a set of *legal-move* and *outcome* functions. Legal-move functions are mappings from dialogues to utterances. Hence, given an (incomplete) dialogue, a legal-move function returns a set of allowed utterances that extend the dialogue. Legal-move functions can also be viewed as functions that specify dialogue constraints. For instance, the *related legal-move* function requires that a latter utterance must be related to some earlier utterance; and the *flat legal-move* function requires that if a sentence has been uttered as the head of a rule, then it is not uttered again as an assumption. Outcome functions are mappings from dialogues to *true/false*. Given a dialogue, an outcome function returns true if a certain property holds within that dialogue. For instance, the *last-word outcome* function returns true if the fictitious proponent agent answers all attacks made by the fictitious opponent agent.

Through dialogues, the participating agents construct a "joint knowledge base" by pooling all information disclosed in the dialogue to form the *ABA framework drawn from a dialogue*, $\mathscr{F}_\delta$. Since a $\mathscr{F}_\delta$ contains all information that the two agents have uttered in the dialogue, it gives the context of examining the acceptability of the claim of the dialogue. Conceptually, a dialogue is "successful" if its claim is acceptable in $\mathscr{F}_\delta$. This soundness result is obtained by mapping the *debate tree* generated from a dialogue to an *abstract dispute tree* [25] that has been developed to prove acceptability results for arguments for various argumentation semantics. This result can be used to prove that certain kinds of these dialogues are successful in resolving conflicts and thus supporting deliberation [**?**].

Some of the earlier study on dialogue systems have categorised dialogues into six types: *persuasion, negotiation, inquiry, deliberation, information-seeking* and *eristics* [116]. It is easy to imagine that each of these types of dialogues has its own goals; and agents participating in different types of dialogue have different interests. Hence different types of dialogues call for different dialogue strategies.

Building upon the aforementioned dialogue protocol, dialogue strategies can be formulated via *strategy-move* functions [**?**, **?**]. These are mappings from dialogues and legal-move functions to utterances. Hence, given a dialogue and a legal-move function, the legal-move function returns a set of utterances that are compatible

[10] Rules, assumptions, and contraries are components of ABA.

with the dialogue protocol; and a strategy-move function selects a subset from these allowed utterances such that utterances within this subset advance the dialogue towards its specific goal.

For instance, in an information-seeking dialogue, where a questioner agent poses a topic and an answerer agent puts forward information that is related to the topic. The behaviours of the questioner and the answerer can be captured in two strategy-move functions: the *pass* and the *non-attack* strategy-move functions, respectively [**?**]. Agents (questioners) that use the pass strategy-move function put forward the claim and no any other utterance in a dialogue; agents (answerers) that use the non-attack strategy-move function only utter rules and assumptions, but not contraries.

Similarly, in an inquiry dialogue, both agents are interested in investigating the acceptability of a given topic. Hence, both agents should be honest and utter all information that each of them knows about the topic. This behaviour can be captured in *truthful* and *thorough* strategy-move functions [**?**], where the truthful strategy-move function selects utterances from one agent's own knowledge base and the thorough strategy-move functions does not select pass if there is any other possible utterances for the agent to make.

In order to support persuasion dialogues, *proponent* and *opponent* strategy-move functions can be used to guarantee that agents are truthful [**?**].

1.3.3 The Argument Web

The plethora of argument visualisation and mapping tools [62] testifies to the enabling function of argumentation-based models for human clarification and understanding, and for promoting rational reasoning and debate. The proliferation of opinion gathering resources and discussion forums on the web, and their lack of support for checking the relevance and rationality of online discussion and debate, has led to increased focus on developing online versions of the aforementioned tools. The advent of such tools in turn raises the possibility of re-use of *ready made* arguments authored online (one of the key issues highlighted by IMPACT is the mining of arguments from online resources). To facilitate both the development of such tools and the reuse of authored arguments, researchers have proposed a need for engineering new systems and standards into the heart of the Internet, to encourage debate, to facilitate good argument, and to promote a new online critical literacy. This is the vision of the *Argument Web* [93]. The Argument Web serves as a common platform that brings together applications in different domains (e.g. broadcasting, mediation, education and healthcare) and interaction styles (e.g. online argument analysis, real-time online debate, blogging). Online infrastructure for argument is combined with software tools that make interacting with the argument web easy and intuitive for various audiences. The infrastructure is built on a putative standard for argument representation, the Argument Interchange Format or *AIF* ([22, 93]). The software tools allow for interactions with the structures represented by the AIF that naturally allow people to express their opinions and link them to those of others, and to use debate as a way of navigating complex issues. The main idea of the Argument Web is visualised in Fig. 1.1. In what follows, we provide a number of examples of specific interactions with the Argument Web, illustrating with prototype tools developed at the School of Computing, University of Dundee.

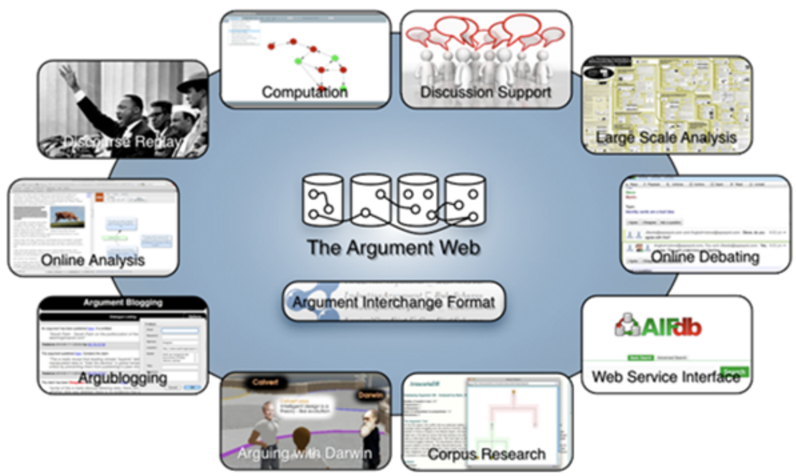

Fig. 1.1 The Argument Web

Arvina

Direct and real-time discussions between two or more people on the web takes place not just via email and instant messaging but also on forums and message boards. These technologies offer only the most basic of structural tools: the discussion is rendered in a linear way and most structure is often brought in by the participants themselves, e.g. by putting "@Chris" in front of their message when they reply to a point made by Chris. The structure of the arguments that are formed in a discussion is thus easily lost.

Our web-based discussion software Arvina [100] allows participants to debate a range of topics in real-time in a way that is structured but at the same time unobtrusive. Arvina uses dialogue protocols to structure the discussion between participants. Such protocols determine which types of moves can be made (e.g. questioning, claiming) and when these moves can be made (e.g. a dialogue starts with a claim, questions can only be moved after a claim has been made). Protocols facilitate a good and rational debate because they, for example, ensure that each participant's opinion is fairly represented and they provide structure to the dialogue itself as well as to the opinions expressed in this dialogue [95]. Fig. 1.2 shows the debate interface. Notice that a (small) part of the Argument Web is displayed as a live discussion map on the right. The argumentative "moves" the user can make in this particular dialogue are represented in the drop-down menu at the bottom.

Fig. 1.2 The Arvina 2 debate interface

In Arvina, reasons for and against opinions are linked to the already available arguments on the Argument Web. Furthermore, Arvina can also use the arguments already on the Argument Web in real-time debate. Arvina takes a multi-agent system populated by agents representing (the arguments of) specific authors who have previously added their opinion to the Argument Web in some way. So, for example, say that Floris has constructed a complex, multi-layered argument using OVA (see below), concerning the use of nuclear weapons. An agent representing Floris can then be added to an Arvina discussion and questioned about these opinions and the agent will answer by giving Floris' opinions. Thus, Arvina cannot just be used to express arguments but also to explore them and to use arguments made by others in one's own reasoning.

OVA

Argument visualisation tools help a user make sense out of a specific complex problem by allowing him to visually structure and analyse arguments. In our opinion, there exists a significant niche market for a lightweight tool which is easily accessible in a browser and makes full use of the functionality provided by the Argument Web. OVA (Online Visualisation of Argument, Fig. 1.3)[11] is a tool for analysing and mapping arguments online. It is similar in principle to other argument analysis tools (being based on Araucaria [94], but is different in that it is accessible from a web browser. This web-based access has allowed for built-in support for

[11] `ova.computing.dundee.ac.uk`

direct analysis of web pages, while also maintaining the ability to analyse standard text files.

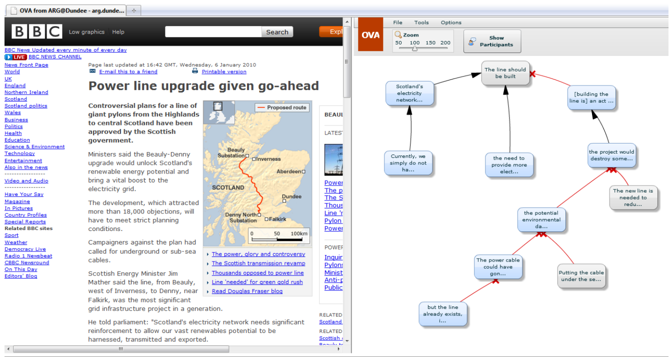

Fig. 1.3 Online Visualisation of Argument (OVA)

OVA is fully linked to the Argument Web and can be used to explore and express arguments in this Argument Web in an intuitive visual manner. One significant advantage that OVA provides over offline packages is that a team of analysts can work together on a problem. Argument analysis is a cognitively intensive and time-consuming task; using OVA, individual analysts can each work on a small part of a complex argument, letting the infrastructure of the Argument Web take care of linking these small subarguments into one complete argument.

Argument Blogging

Our third and final example of how argumentation technologies based on the Argument Web can facilitate online debate concerns the popular activity of blogging. As with message boards, the current structure of the Internet only allows for simple dialogue and argument structures: if one wants to reply to an opinion presented somewhere on the web in one's blog, for example, one can provide a simple hyperlink to the article in which this opinion is expressed. The resulting structure of supporting and competing opinions is then easily lost. Furthermore, because each new claim is expressed on its own page (i.e. someone's blog page) and there is no overview of the dialogue between the various authors and bloggers.

In order to improve rational debate that uses the popular blogging format, we have built a very simple tool for Argument Blogging [119] which allows opinions in blogs and other web pages to be easily linked using the underlying infrastructure of the Argument Web. The tool is essentially an addition to the context-menu in a web browser (the "right-click" menu): when selecting some text on a webpage, the user can opt to either agree or disagree with this text and type in their reasons for (dis)agreement. These reasons are then automatically posted to the user's blog, with a link (URI) to the original text. More importantly, the "agree" or "disagree" argument moves and their resulting claims are all aggregated on the Argument Web. A discussion, which may be the result of multiple subsequent uses of the tool, can then be explored using any other tool for the Argument Web, such as Arvina or OVA.

Linking Computational Models of Argument to Human Authored Arguments

Formal models of argumentation enable the structuring of individual arguments and the dialogical exchange of argument in offline and online tools for argumentation-based human reasoning and debate. Thus far, there has been little work on organising human authored arguments into Dung argumentation frameworks, and evaluating the status of these under Dung's various semantics (see Section 1.1.2). The provision of this evaluative functionality would: 1) ensure that the assessment of arguments is formally and rationally grounded; 2) enable humans to track the status of arguments so that they can be guided in which arguments to respond to; 3) enable 'mixed' argumentation integrating both machine and human authored arguments.

We briefly report on recent work aiming at providing this functionality. Earlier we referred to the Argument Interchange Format (AIF) ([22, 93]) that has been proposed as a standardised format for representation of argumentation knowledge. The idea is that the AIF can serve as a common representation language for human authored arguments and arguments constructed in logic, so that (for example) human authored arguments can be translated to a formal logic representation for evaluation under Dungs semantics. This idea is explored in [14], in which two-way translations between the AIF and the recent $ASPIC^+$ framework [89, 77] are defined. $ASPIC^+$ is a general framework that provides a structured (rather than fully abstract) account of argumentation. The idea is that one can define a range of logic-based instantiations of this structured framework such that the defined arguments and their defeats (attacks that succeed with respect to preferences over arguments) can be evaluated under Dung's semantics, while ensuring that rationality postulates for argumentation [19] are satisfied. One can then take AIF representations of arguments and their interactions defined in the above mentioned tools, and translate these to instantiations of the $ASPIC^+$ framework, so enabling evaluation under Dung's semantics. This is explored in [14], in which arguments and their interactions authored in the *Rationale* tool [10] are translated to the AIF and then to $ASPIC^+$ representations.

1.3.4 Argumentation and the Social Web

In the Social Web, users connect with each other and share knowledge and experiences of all types, in interactions that often resemble debates with exchange of arguments (e.g. in comments on blogs). Nevertheless, the argumentative structure is implicit [98], arguments need to be inferred [108], debates are unstructured, often chaotic [64], not to mention the disruption caused by the Trolls and their inflammatory, extraneous, or off-topic participation [110].

Whereas the use of argumentation in the Social Web context has been advocated by many authors [109, 110, 98, 64, 108] as a channel by means of which argumentation can inform human reasoning, the realisation of such a vision is yet to come.

Most existing work considering online systems and argumentation (some of which discussed is elsewhere in this Chapter) focuses on extracting argumentation frameworks, of one form or another, manually or semi-automatically from user exchanges, e.g. through the use of argument schemes as a way to understand the contributions in these exchanges [58], or by mapping these contributions onto the AIF format, again using argument schemes as well as semantic web technology for editing and querying arguments [93]. These works implicitly assume that the extraction of argumentation frameworks is down to "argumentation engineers", fluent in (one form or another of) computational argumentation. This prevents these systems to scale and be widely adopted in the Social Web.

Recently, some steps that do not assume the existence of such argumentation engineers have been taken, two of which we outline next, one using Abstract Argumentation and the other Assumption Based Argumentation.

1.3.4.1 Social Abstract Argumentation

In [64], Leite and Martins introduce the notion of *Social Abstract Argumentation Frameworks*, an extension of Dung's AAF with the possibility to associate votes to arguments.

Social Abstract Argumentation Frameworks are meant to provide formal support to self-managing online debating systems capable of accommodating two archetypal levels of participation. On the one hand, experts, or enthusiasts, will be provided with simple mechanisms to specify their arguments and also a way to specify which arguments attack which other arguments. To promote participation, arguments can be anything such as a textual description of the argument, a link to some source, a picture, or any other piece of information these users deem fit. On the other hand, less expert users who prefer to take a more observational role will be provided with simple mechanisms to vote on individual arguments, and even on the specified attacks. The system will then be able to autonomously determine outcomes of debates – the social value of arguments – taking into account the structure of the argumentation graph consisting of arguments and attacks, and the crowd's opinion expressed by the votes. These will be fed to a GUI, which will display arguments and attacks with shades or sizes proportional to their strengths, while adapting as new arguments, attacks and votes are added, thus enabling users to observe the current state of the debate.

In [64], the authors define a class of semantics for Social Abstract Argumentation Frameworks where the social value assigned to arguments goes beyond the usual accepted/defeated and can take values from any arbitrary set of values. Some of the proposed semantics exhibit several formal properties which can be mapped to desirable features of the online debating system, related to democracy, universality, etc. According to Leite and Martins, the use of abstract argumentation allow great flexibility in the process of specifying arguments, thus fostering participation by allowing users with different levels of expertise to be able to easily express their arguments.

We illustrate some possible novel uses of argumentation in a Social Web context:

Participatory journalism

Let us consider the following (fictitious) scenario. User Bob is reading an online newspaper. He just finished reading a controversial article on *Do Androids Dream of Electric Sheep?* and he wants to share his thoughts on the matter. But there are already 1,357 user comments! The two comments at the top of the page seem quite interesting. Next to the first comment Bob reads 45 people like this. 32 people like the second one. Then there are some adverts. After that, there are a couple of recently added comments, followed by an older and quite long thread of insults, directed to readers, androids, and sheep alike. Now Bob's problems are:

- gosh, there is so much noise in this discussion - what do people think about this article? does anyone feel like me about it? I don't want to read 1,364 user comments[12]
- if I write something that has already been said, are people going to insult me?
- I don't know any of these people writing their comments here – is there anyone who knows what he is talking about?
- what did I want to write? I forgot
- what was the article about?

If we think of it, the management of debates in current Social Web sites is very primitive. There are no solutions that can solve Bob's problems. The more people give their contribution, the less their contribution is usable, because there is too much noise. This is because these technologies do not have debate-oriented concepts. Argumentation can provide these concepts.

We envisage the possibility of a new participatory journalism web site (let us call it `ArguingtonPost.com`) empowered with argumentation and voting technologies. It provides many innovative debate-oriented features such as: visualizing comments in a more usable way, e.g. by clustering comments that agree with each other; maintaining collateral user information, such as the user's authority on specific subjects, as emerged from previous discussion, or its positive/negative contribution to discussion; filtering out comments posted by trolls and grievers; promoting connections between users who agree on similar positions.:

Sentiment-aware search engines

Modern search engines represent, for a large share of Internauts, the "Portal" to the World Wide Web. If you want to know what is a "gridiron," or how "George Benson" looks like, or where to go "out for dinner in Kowloon", or "how to prepare

[12] Meanwhile, some more insults appeared, which increased the comments counter.

tiramisù" you just type a couple of keywords in Google or Yahoo and get the answer. Well, let's say you get a number of possible links, ranked by very smart algorithms, and a bunch of related adverts. In many cases, you are lucky and the first or second link is what you need.

This was true until just recently. Now the way people access the Web is changing. Instead of typing your queries in a search engine, you can change your status in Facebook or Twitter, saying for example you're preparing a tiramisù for your darling, and some of your friends will probably give you tips and links with tested receipts.[13] A possible reason of this change in the Internauts' life style is that your Facebook friends will actually give you better information, and pester less with useless spam. Indeed, search engines are interested in opinion mining and sentiment analysis [49, 85] and we expect this hot research area to produce very interesting results in the near future.

We envisage a sentiment-aware search engine (let us call it `Arguugle.com`) that mines large online discussion boards that use technologies such as the aforementioned `ArguingtonPost.com`. In this way, `Arguugle.com` can offer some innovative features, including advanced clustering of result based on user agreement/sentiment, and guessing of user intention and display of additional, not-asked-for information such as positions "in favour" or "against", tips and advice.

Advanced ranking.

Suppose that I never read any novel by Stephen King. I want to start with a good one. I type "Wiki Stephen King's novels" on Google. The first hit is Stephen King's Wikipedia main article. The second one is Wikipedia's article about the novel *It*. The third hit is Stephen King's bibliography's Wikipedia main article (obviously what I was looking for). The fourth hit is Listverse's "Top 15 Stephen King Books."

Listverse is a "Top 10 List" web site. That particular ranking is made by a user called Mon. At the time of writing, that ranking has 535 comments of users who agree or disagree with Mon's list. The Web is full of web sites like Listverse: Rankopedia, Squidoo, lists by newspapers such as the Guardian, the Times or USA Today, bookseller lists like Amazon and Barnes & Noble, etc. Ranking and recommendation are everywhere, because they can help us every time we must make a decision about things that require expertise we don't have. Where shall I stop in my Andalucia tour? Which optic is best for my camera? Who's the best catcher of all times?

Recommendation web sites can be very simple: just a numerical ranking, as a result of voting. Or they can require some expert to write their opinion and people to comment. Some popular recommendation services for trip organization, typically

[13] A 2010 survey illustrates Facebook overtaking Google's popularity among US Internet users. See "Facebook becomes bigger hit than Google" by By Chris Nuttall and David Gelles on Financial Times, online March 17, 2010 `www.ft.com/cms/s/2/67e89ae8-30f7-11df-b057-00144feabdc0.html#axzz1MSvZe0pb`. Recently Facebook is investing on a "social web search" project in order to better exploit its social data. See "Facebook Delves Deeper Into Search" By Douglas MacMillan and Brad Stone on Bloomberg Business Week, online March 29, 2012 `www.businessweek.com/articles/2012-03-28/facebook-delves-deeper-into-search`.

associated with online hotel booking services, divide comments into positive and negative ones. That helps. But in general, as a lazy user, I don't want to read too much text, and at the same time I am not impressed by crude rankings because I want to know *what* people give value to when they say "this hotel if fabulous" or "that book is boring".

We envisage an argumentation-empowered recommendation web site (let us call it `Argubest.com`) that uses argumentation technologies and is able to:

- use numerical rating together with user comments and relations between comments when computing the ranking, thus providing a very convincing ranking
- organize feedback and opinions in a simple and intuitive way for the user to browse them
- understand which comments seem to be misleading or of little use, and filter them out
- understand which users seem to be more reliable and give more importance to their ratings

1.3.4.2 Bottom-Up Argumentation

In [108], Toni and Torroni propose the use of Assumption-Based Argumentation to assess the dialectical validity of the positions debated in, or emerging from the exchanges in online social platforms.

They envisage a system where active participants in the exchange are annotating the exchanges, where annotations indicate that pieces of text in natural language are either comments or opinions, and links can be drawn to indicate source, support or objection. Users will add comments, opinions and links dynamically, in the same way as exchanges grow over time in existing online systems. These annotations are then mapped to an existing computational argumentation framework, Assumption-Based Argumentation (ABA), paving the way to the automatic computation of the dialectical validity of comments, opinions, and links, and thus topics that these encompass.

According to Toni and Torroni, the use of ABA as the underlying computational argumentation framework is justified by the fact that it is equipped with a variety of well-defined semantics and computational counterparts for assessing dialectical validity, its ability to distinguish arguments, support as well as attack amongst them, and its capability of dealing with defeasibility of information, important as the system evolves over time.

Whereas both Social Abstract Argumentation and ABA based Bottom-Up Argumentation both share the view that users, instead of specialised "argumentation engineers", share the burden of defining the structure of the argumentation framework, some features set them apart. In Social Abstract Argumentation, users are allowed to vote on arguments and attack relations, and the votes dynamically reflect on the gradual value of arguments – implementing a more subjective view on argumentation which is perhaps closer to real interactions in the Social Web where consensus hardly ever exists. In ABA based Bottom-Up Argumentation, there is no counterpart to voting and the underlying semantics sticks to the classical accepted/defeated assignment. Then, in contrast to Social Abstract Argumentation, ABA based Bottom-Up Argumentation permits the specification of a support relation which is a common feature of most interactions in the Social Web. Perhaps a combination of both is the best approach to better reflect what goes on in the Social Web, adopting the votes and gradual values from Social Abstract Argumentation and the support relation from ABA based Bottom-Up Argumentation.

1.3.4.3 Discussion: Open Issues and Challenges

We discuss here several challenges that lay ahead for a full integration of argumentation in a Social Web context.

Firstly, the use of Argumentation in Social Computing requires the development of a suitable underlying knowledge representation framework that accommodates all the information provided by the users, together with a semantics that combines an argumentation framework with the community feedback to assign a value to each argument. We need to understand and formalize new concepts such as "social support", "social acceptability", to describe the positions of the community with respect to the matter under discussion. Such semantics should exhibit several desirable properties, to ensure acceptance of the outcomes and promote appropriate user behaviour.

There are many suitable candidates for the basic argumentation framework: Abstract Argumentation Frameworks [29], Value-based Argumentation Frameworks [7], Assumption-based Argumentation Frameworks [26], Meta-level Argumentation Frameworks [75]. Recently, Leite and Martins [64] introduced Social Abstract Argumentation Framework which allows to attach votes to abstract arguments and exhibits several desirable semantical properties for using it in Social Computing.

Secondly, successful Social Computing services are based on few mechanisms, which are already known to the user, or easy to be learned. In many applications, information and social exchange has an entertainment component. The use of argumentation in social computing introduces an additional level of structure in interactions which will bring additional challenges in the development of interfaces. This new class of interfaces should be simple enough to be engaging, but at the same time allowing for richer interactions, accommodating the participation of users with various degrees of expertise and motivation.

The interface must provide, for all kinds of users, the right level of abstraction that allows them to interact at the desired level of detail by adding content, identifying relations between claims, navigating through the debate, etc., or simply by voting. As a debate proceeds, the interface will perhaps resort to colors, fonts, geometries or other visual artifacts to highlight a prevailing opinion, and emphasize agreements, supporting arguments, attacks and contradictions. Existing visualization tools [62] could be used to enhance clarity of presentation and promote user acceptance.

Thirdly, a key challenge is the development of efficient algorithms that can effectively support argumentation and voting together at run time and at a large scale (comments, users). Such algorithms will have to rapidly propagate the effect of changes in a debate, be it a new argument or simply a new vote.

Finally, automated text extraction is one of the most challenging problems in any application that involves knowledge intensive interaction between man and machine. Techniques that automatically identify claims from human-generated text would enable automating tasks such as establishing relations between claims, checking for consistency, etc. Recent advances in automated text extraction and, specifically, on Web dispute identification [37], lead us to believe that soon the technology will be ripe to identify claims in discussion forums effectively and automatically, or at least semi-automatically (e.g. with the help of the social community). An increase in the efficiency of automated text extraction and claim identification will be accompanied with a significant increase in the potential for use of argumentation in the Social Web.

1.4 Benchmark Libraries for Argumentation

For formal models of argumentation to inform human and machine reasoning, argumentation needs to be supported by computational systems and tools. The argumentation community has been very active in the last decade in delivering argumentation engines. Several dedicated engines have been released, such as, for instance, DeLP[14] for the argumentation framework of [45], the CaSAPI system[15] for the Assumption-Based Argumentation (ABA) framework of [17, 25, 27, 26], CARNEADES[16] for the argumentation framework of [54], the ASPIC system[17], for the argumentation framework of [89], as well as an increasing number of implementations for computing extensions in abstract argumentation [29]. Well-known representatives of this latter class of systems are Verheij's system[18] [114], ArguLab[19] [86], and ASPARTIX[20] [35].

While the former two are based on tailored algorithms for abstract argumentation, ASPARTIX follows a reduction approach where the actual computation is delegated to an ASP-engine[21]. A number of other approaches using ASP have also been proposed (see [107] for a survey). A similar approach has been followed in [15], suggesting to use CSP solvers for the main computations. Another option would be

[14] `http://lidia.cs.uns.edu.ar/delp_client/`

[15] `http://www.doc.ic.ac.uk/~ft/CaSAPI/`

[16] `http://carneades.berlios.de/`

[17] `http://www.arg.dundee.ac.uk/toast/`

[18] `http://www.ai.rug.nl/~verheij/comparg/`

[19] `http://heen.webfactional.com/`

[20] `http://rull.dbai.tuwien.ac.at:8080/ASPARTIX`

[21] Answer-Set Programming (ASP) [84] is a declarative programming paradigm which allows for succinct representation of combinatorial problems.

to employ SAT-solvers, as discussed by Besnard and Doutre [12], or QSAT-solvers, as discussed by Egly and Woltran [36] (however implemented systems of these two kinds are not available yet). Recent work demonstrates that other methods are also applicable to abstract argumentation, in particular dynamic algorithms based on tree decompositions [34] or computations based on backdoor sets [33]. All of the mentioned systems or proposals cover a certain range of abstract argumentation semantics (see e.g. [6] for an overview), but nearly all of them include Dung's standard semantics, such as preferred, stable, or complete extensions [29].

Considering the number of proposed argumentation systems, we believe that a benchmark library is indispensable for a systematic comparison and evaluation thereof, with an eye towards application scenarios and deployment in applications. We shall highlight here some main requirements for such a library, have a look at similar such collections in other areas, and raise some questions which the argumentation community should consider and agree upon. We will focus on abstract argumentation systems and consider the following issues:

- How to compare the performance of the different systems for abstract argumentation?
- How to verify the correctness of the systems?
- To which level of problem size do current approaches scale well?
- How can data between different applications and solvers be exchanged?
- How can we - in the long term - measure the progress the community makes in terms of practical systems?

We will advocate the importance of a benchmark library as a way to address these issues. We will also discuss general issues like suitable input formats. Taking the wide variety of extensions of abstract argumentation into account, such a format should be extendable in the sense that, for example, value-based argumentation frameworks (VAFs) [7] and argumentation frameworks with recursive attacks (AFRAs) [5, 74] can be captured as well.

It is worth mentioning at this point that abstract argumentation itself is not the only framework available and may not be suitable for all applications (see e.g. [19]) and abstract argumentation systems are only some of the available engines, as our earlier discussion shows. Nonetheless, efficient systems for abstract argumentation deserve attention as they are an important step towards handling problems of real-world size in order to prolong the success-story of argumentation within the AI community.

Finally, we will raise concrete questions about how a benchmark library for argumentation should be set up and also have a look at how other communities dealt with this kind of service.

1.4.1 The Value of a Benchmark Library

The following thesis was proposed by Toby Walsh in his talk at the 2009 AAAI Spring Symposium[22].

Every mature field has a benchmark library.

We would like to subscribe to this thesis and paraphrase some of the general benefits of a benchmark library as pointed out by Walsh.

With a growing and well-maintained benchmark library for argumentation, researchers can test their ideas and concepts on instances from a wide range of applications. If the library includes instances of different size, from a few dozens to thousands of arguments, one can use it to evaluate how well an algorithm or reasoning method scales, to which kind of instances it applies best, and to which kind of instances it does not. In order to establish and maintain a useful library, the research community should therefore be encouraged to contribute benchmark instances. Instances should be from various categories, including random instances, hand-crafted instances, and instances that arise from real-world instantiations of argumentation.

A benchmark library will bring various benefits to the field of argumentation as it will support the implementation of new theoretical ideas, as well as their testing and comparison with the state of the art. It will also reward efforts put into the engineering part of the implementation, and so support the combination of theoretical and practical contributions. A benchmark library will highlight some low level aspects that are easily overlooked by a purely abstract theoretical treatment.

For instance, research on propositional satisfiability (SAT) has enormously benefited from a large and diverse benchmark library (see [59] and respectively `http://www.satlib.org/`). By means of a benchmark library one can witness the progress over the years. For SAT, the size of solvable instances increased by an order of magnitude every ten years since the 1980s (see, e.g. [11] or `http://www.satcompetition.org` for the more recent progress).

A well-maintained benchmark library is a necessary prerequisite for a solver competition. There can be a benchmark library without a competition, but no competition without a benchmark library. Maybe in a couple of years argumentation will be ready for such a competition.

1.4.2 Towards a benchmark library

Following Toby Walsh a benchmark library should be located on the web and easy to find. We would suggest to use `http://www.arglib.org` (following the naming convention from other related areas, e.g. `http://www.csplib.org` and `http://www.satlib.org`) and we have already reserved it for this purpose.

[22] Benchmarking of Qualitative Spatial and Temporal Reasoning Systems, Stanford University, CA, USA, March 23-25, 2009

In what follows, we thus use `arglib` as a shorthand for the library we have in mind.

To set up `arglib`, an important issue is to find an appropriate format to represent instances. The following points can be made:

- It has to be decided whether a hi-tech format like XML or a lo-tech format like DIMACS, which is successfully used in SATLib [59], shall be used.
- The format should be non-proprietary and widely accepted by the community.
- On the one hand we would like a simple representation of abstract argumentation frameworks. On the other hand, the format should be able to capture extensions of Dung's abstract frameworks like the aforementioned VAFs, AFRAs, and many others. In addition, the format should allow to represent information about the internal structure of arguments in case they are obtained from an instantiation process. As the argumentation community is widespread and frequently comes up with new formal systems, it is very unlikely that one can provide a format capturing all relevant ideas. Hence we seek for a format that is both *simple* and *easily extendable*.

  A potential role model could be the *UAI file format* used for benchmarking probabilistic reasoning problems.[23] For probabilistic reasoning, one takes as input a graphical model of a probability distribution which consists of a graph whose nodes are annotated with numerical values or tables. The UAI file format uses for that purpose a simple ASCII text file. The first part of the file represents the *graph structure* of the graphical model, the second part represents the *annotations*. A similar approach might be useful for argumentation, where one could use the first part to represent the basic attack relation, the second part to represent additional information such as preferences, weights, etc. The first part would remain the same for exchanging data for a wide range of argumentation systems, whereas the second part could provide some flexibility for special application or extensions of basic abstract argumentation frameworks.

We believe that the existing argument interchange format [92], AIF for short, is not well suited for `arglib`. In particular, this format was introduced with a different motivation, namely to have a common ontology supporting interchange between different argumentation approaches and systems. Thus its facilities go far beyond the purely abstract formalisms we consider here. Although AIF provides a rich framework to specify graphs (via its so-called upper ontology), we believe that for the purpose of a benchmark library for abstract argumentation systems a simple format is the better choice. Once translations between a simple format and AIF are established, there might be also the opportunity to extract benchmarks directly from AIF specifications.

[23] `www.cs.huji.ac.il/project/UAI10/fileFormat.php`

1.4.3 Instances

It is obvious that `arglib` should offer a broad range of instances.

- There should be small (maybe hand-crafted) instances as well as huge instances. This allows to test and compare how different solvers scale.
- A simple way of generating instances is to use a random generator. Such random instances have the advantage that one can produce a wide range of instances with increasing size and gradual changes in density. However, random instances have the disadvantage that they lack the typical structure that is present in real-world instances, hence using them alone for measuring the performance of a solver can produce misleading results, and optimizing a solver solely on random instances is not useful for its performance in practise.
- Real-world instances should be obtained from various applications and different kind of instantiations to avoid that `arglib` becomes biased.
- For solved instances, the solutions should be available as well. This would allow to empirically verify new solvers.

We conclude by mentioning that a successful library needs the support of the community. Who should maintain `arglib`: a consortium, a research group, or even just a single person? To build a representative library it is important that researchers submit benchmarks. So inevitably the question arises as to how to motivate the community to submit their examples to `arglib`?

There are several related research areas close to argumentation that already have widely accepted benchmark libraries. Hence it might be a good idea to learn from them. A joint workshop with organisers from other areas such as SAT, CSP, or ASP [24] might be a starting point.

1.5 Conclusions

We have provided an overview of a number of approaches relying upon argumentation to either support humans or machines towards reaching agreement. Examples of argumentation-augmented machine reasoning include methods for machine learning and trust computing. Examples of argumentation in support of human reasoning include several forms of (individual and collaborative) decision-making and methods in the context of the Web and Social Networks. We have also discussed some open issues, in the context of the individual approaches surveyed as well as in general, for argumentation to strengthen its potential and further demonstrate the added value it brings to applications. Concretely, we have identified the need for benchmark libraries as an important open challenge.

The approaches we have overviewed witness the added value of argumentation in a number of settings. Recent work in cognitive science and psychology [72] gives an argumentation-based account of how human capacity to reason evolved. This theory further suggests the use of argumentation in supporting humans to arrive at better outcomes when engaged in the interactive process of arriving at agreements.

Acknowledgements Many thanks to Jordi Sabater-Mir and Vicente Botti for useful suggestions and comments on an earlier version of this chapter.

References

1. Amgoud, L.: Modeling dialogues using argumentation. In: ICMAS '00: Proceedings of the Fourth International Conference on MultiAgent Systems, p. 31. IEEE Computer Society, Washington, DC, USA (2000)
2. Amgoud, L., Cayrol, C.: A reasoning model based on the production of acceptable arguments. Annals of Mathematics and Artificial Intelligence **34**(1-3), 197–215 (2002)
3. Amgoud, L., Cayrol, C., Lagasquie-Schiex, M., Livet, P.: On bipolarity in argumentation frameworks. International Journal of Intelligent Systems **23**(10), 1062–1093 (2008)
4. Atkinson, K.: What should we do?:computational representation of persuasive argument in practical reasoning. Ph.D. thesis, Department of Computer Science, University of Liverpool, UK (2005)
5. Baroni, P., Cerutti, F., Giacomin, M., Guida, G.: Afra: Argumentation framework with recursive attacks. International Journal of Approximate Reasoning **52**(1), 19–37 (2011)
6. Baroni, P., Giacomin, M.: Semantics of abstract argument systems. In: I. Rahwan, G. Simari (eds.) Argumentation in Artificial Intelligence, pp. 25–44. Springer (2009)
7. Bench-Capon, T.J.M.: Value-based argumentation frameworks. In: Proceedings of the 10th International Workshop on Non-Monotonic Reasoning (NMR'02), pp. 443–454 (2002)
8. Bench-Capon, T.J.M.: Persuasion in practical argument using value-based argumentation frameworks. Journal of Logic and Computation **13**(3), 429–448 (2003)
9. Bench-Capon, T.J.M., Dunne, P.E.: Argumentation in artificial intelligence. Artificial Intelligence **171**, 10–15 (2007)
10. Berg, T., van Gelder, T., Patterson, F., Teppema, S.: Critical Thinking: Reasoning and Communicating with Rationale. Amsterdam: Pearson Education Benelux (2009)
11. Berre, D.L., Simon, L.: Preface. Journal on Satisfiability, Boolean Modeling and Computation **2**(1-4) (2006)
12. Besnard, P., Doutre, S.: Checking the acceptability of a set of arguments. In: Proceedings of the 10th International Workshop on Non-Monotonic Reasoning (NMR'02), pp. 59–64 (2004)
13. Besnard, P., Hunter, A.: A logic-based theory of deductive arguments. Artificial Intelligence **128**(1–2), 203–235 (2000)
14. Bex, F., Modgil, S., Prakken, H., Reed, C.: On logical reifications of the argument interchange format. Journal of Logic and Computation **To appear** (2012)
15. Bistarelli, S., Santini, F.: A common computational framework for semiring-based argumentation systems. In: H. Coelho, R. Studer, M. Wooldridge (eds.) Proceedings of the 19th European Conference on Artificial Intelligence (ECAI'10), *Frontiers in Artificial Intelligence and Applications*, vol. 215, pp. 131–136. IOS Press (2010)
16. Bochman, A.: Collective argumentation and disjunctive programming. Journal of Logic and Computation **13 (3)**, 405–428 (2003)
17. Bondarenko, A., Dung, P., Kowalski, R., Toni, F.: An abstract, argumentation-theoretic approach to default reasoning. Artificial Intelligence **93**(1-2), 63–101 (1997)
18. Bratko, I., Žabkar, J., Možina, M.: Argument based machine learning. In: I. Rahwan, G. Simari (eds.) Argumentation in Artificial Intelligence, pp. 463–482. Springer Verlag (2009)
19. Caminada, M., Amgoud, L.: On the evaluation of argumentation formalisms. Artificial Intelligence **171**(5-6), 286–310 (2007)

20. Cayrol, C., Lagasquie-Schiex, M.C.: Graduality in argumentation. J. of Art. Int. Res. **23**, 245–297 (2005)
21. Chesñevar, C., Maguitman, A., González, M.P.: Empowering recommendation technologies through argumentation. In: I. Rahwan, G. Simari (eds.) Argumentation in Artificial Intelligence, pp. 403–422. Springer (2009)
22. Chesñevar, C., 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**, 293–316 (2006)
23. Clark, P., Boswell, R.: Rule induction with CN2: Some recent improvements. In: Machine Learning - Proceeding of the Fifth Europen Conference (EWSL-91), pp. 151–163. Berlin (1991)
24. Denecker, M., Vennekens, J., Bond, S., Gebser, M., Truszczynski, M.: The second answer set programming competition. In: E. Erdem, F. Lin, T. Schaub (eds.) Proceedings of the 10th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR 2009), *LNCS*, vol. 5753, pp. 637–654. Springer (2009)
25. Dung, P., Kowalski, R., Toni, F.: Dialectic proof procedures for assumption-based, admissible argumentation. Artificial Intelligence **170**, 114–159 (2006)
26. Dung, P., Kowalski, R., Toni, F.: Assumption-based argumentation. In: I. Rahwan, G. Simari (eds.) Argumentation in AI, pp. 199–218. Springer (2009)
27. Dung, P., Mancarella, P., Toni, F.: Computing ideal sceptical argumentation. Artificial Intelligence, Special Issue on Argumentation in Artificial Intelligence **171**(10–15), 642–674 (2007)
28. Dung, P.M.: An argumentation semantics for logic programming with explicit negation. In: Proceedings of the Tenth Logic Programming Conference, pp. 616–630. MIT Press, Cambridge, MA (1993)
29. Dung, P.M.: On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. Artificial Intelligence **77**(2), 321–358 (1995)
30. Dung, P.M., T., P.M., Hung, N.D.: Argument-based decision making and negotiation in e-business: Contracting a land lease for a computer assembly plant. In: M. Fisher, F. Sadri, M. Thielscher (eds.) Proceedings of the 9th International Workshop on Computational Logic in Multi-Agent Systems (CLIMA IX), *Lecture Notes in Computer Science*, vol. 5405, pp. 154–172. Springer (2009)
31. Dung, P.M., Thang, P.M., Toni, F.: Towards argumentation-based contract negotiation. In: Proc. COMMA'08. IOS Press (2008)
32. Dunne, P.E., Hunter, T., McBurney, P., Parsons, S., Wooldridge, M.: Weighted argument systems: Basic definitions, algorithms, and complexity results. Artificial Intelligence **175**, 457–486 (2011)
33. Dvořák, W., Ordyniak, S., Szeider, S.: Augmenting tractable fragments of abstract argumentation. Artificial Intelligence (2012). DOI 10.1016/j.artint.2012.03.002. In press (Available online 8 March 2012)
34. Dvořák, W., Pichler, R., Woltran, S.: Towards fixed-parameter tractable algorithms for abstract argumentation. Artificial Intelligence **186**(0), 1 – 37 (2012)
35. Egly, U., Gaggl, S.A., Woltran, S.: Answer-set programming encodings for argumentation frameworks. Argument and Computation **1**(2), 147–177 (2010)
36. Egly, U., Woltran, S.: Reasoning in argumentation frameworks using quantified boolean formulas. In: P.E. Dunne, T.J.M. Bench-Capon (eds.) Proceedings of the 1st Conference on Computational Models of Argument (COMMA'06), *Frontiers in Artificial Intelligence and Applications*, vol. 144, pp. 133–144. IOS Press (2006)
37. Ennals, R., Trushkowsky, B., Agosta, J.M.: Highlighting disputed claims on the web. In: Proc. 19th WWW, pp. 341–350. ACM (2010)
38. Falappa, M., García, A., Kern-Isberner, G., Simari, G.: On the evolving relation between belief revision and argumentation. Knowledge Eng. Review **26**(1), 35–43 (2011)

39. Falappa, M., Kern-Isberner, G., Simari, G.R.: Argumentation in Artificial Intelligence, chap. Belief Revision and Argumentation Theory, pp. 341–360. I. Rahwan, G. R. Simari (eds.), Springer (2009)
40. Fan, X., Toni, F.: Assumption-based argumentation dialogues. In: Proc. IJCAI (2011)
41. Feigenbaum, E.A.: Some challenges and grand challenges for computational intelligence. Source Journal of the ACM **50**(1), 32–40 (2003)
42. Ferretti, E., Errecalde, M., García, A., Simari, G.R.: Decision rules and arguments in defeasible decision making. In: P.B. et al (ed.) Proc. 2nd Int. Conf. on Computational Models of Arguments (COMMA), *Frontiers in Artificial Intelligence and Applications*, vol. 172, pp. 171–182. IOS Press (2008)
43. Fox, J., Glasspool, D., Grecu, D., Modgil, S., South, M., Patkar, V.: Argumentation-based inference and decision making–a medical perspective. IEEE Intelligent Systems **22**(6), 34–41 (2007). DOI 10.1109/MIS.2007.102. URL `http://dx.doi.org/10.1109/MIS.2007.102`
44. García, A., Rotstein, N., Tucat, M., Simari, G.R.: An argumentative reasoning service for deliberative agents. In: LNAI 4798. KSEM 2007, pp. 128–139. Springer (2007)
45. García, A., Simari, G.: Defeasible logic programming: an argumentative approach. Theory and Practice of Logic Programming **4**(1), 95–138 (2004)
46. García, A.J., Rotstein, N., Chesñevar, C., Simari, G.R.: Explaining why something is warranted in defeasible logic programming. In: T.R.B. et al (ed.) ExaCt, pp. 25–36 (2009)
47. García, D., Gottifredi, S., Krümpelmann, P., Thimm, M., Kern-Isberner, G., Falappa, M., García, A.: On influence and contractions in defeasible logic programming. In: J.P. Delgrande, W. Faber (eds.) LPNMR, *Lecture Notes in Computer Science*, vol. 6645, pp. 199–204. Springer (2011)
48. Glasspool, D., Oettinger, A., Smith-Spark, J., Castillo, F., Monaghan, V., Fox, J.: Supporting medical planning by mitigating cognitive load. Methods of Information in Medicine **46**, 636–640 (2007)
49. Godbole, N., Srinivasaiah, M., Skiena, S.: Large-scale sentiment analysis for news and blogs. In: Proceedings of the International Conference on Weblogs and Social Media (ICWSM) (2007)
50. González, M., Chesñevar, C., Pinwart, N., Gomez Lucero, M.: Developing argument assistant systems from usability viewpoint. In: Proc. Int. Conference on Knowledge Management and Information Sharing, pp. 157–163. INSTICC (2010)
51. González, M., Lorés, J., Granollers, T.: Enhancing usability testing through datamining techniques: A novel approach to detecting usability problem patterns for a context of use. Information & Software Technology **50**(6), 547–568 (2008)
52. González, M.P., Gottifredi, S., García, A.J., Simari, G.R.: Towards argument representational tools for hybrid argumentation systems. In: G. Salvendy, M.J. Smith (eds.) HCI (12), *Lecture Notes in Computer Science*, vol. 6772, pp. 236–245. Springer (2011)
53. Gordon, T.F.: The Pleadings Game. An Artificial Intelligence Model of Procedural Justice. Kluwer Academic Publishers, Dordrecht/Boston/London (1995)
54. Gordon, T.F., Walton, D.: The Carneades argumentation framework - using presumptions and exceptions to model critical questions. In: P.E. Dunne, T.J.M. Bench-Capon (eds.) Proceedings of the 1st Conference on Computational Models of Argument (COMMA'06), *Frontiers in Artificial Intelligence and Applications*, vol. 144, pp. 195–207. IOS Press (2006)
55. Governatori, G., Maher, M.J.: An argumentation-theoretic characterization of defeasible logic. In: Proceedings of the Fourteenth European Conference on Artificial Intelligence, pp. 469–473 (2000)
56. Groznik, V., Guid, M., Sadikov, A., Možina, M., Georgijev, D., Kragelj, V., Ribarič, S., Pirtošek, Z., Bratko, I.: Elicitation of neurological knowledge with ABML. In: Proc. of the 13th Conference on Artificial Intelligence in Medicine (AIME'11), July 2-6, Bled, Slovenia (2011)
57. Hamblin, C.L.: Mathematical models of dialogue. Theoria **37**, 130–155 (1971)

58. Heras, S., Atkinson, K., Botti, V.J., Grasso, F., J., V., McBurney, P.: How argumentation can enhance dialogues in social networks. In: P. Baroni, F. Cerutti, M. Giacomin, G.R. Simari (eds.) Computational Models of Argument: Proceedings of COMMA 2010, Desenzano del Garda, Italy, September 8-10, 2010, *Frontiers in Artificial Intelligence and Applications*, vol. 216, pp. 267–274. IOS Press (2010)
59. Hoos, H.H., Stützle, T.: Satlib: An online resource for research on SAT. In: Proceedings SAT 2000, pp. 283–292. IOS Press (2000)
60. Hussain, A., Toni, F.: On the benefits of argumentation for negotiation - preliminary version. In: Proceedings of 6th European Workshop on Multi-Agent Systems (EUMAS-2008) (2008)
61. Jennings, N.R., Faratin, P., Lomuscio, A.R., Parsons, S., Sierra, C., Wooldridge, M.: Automated negotiation: Prospects, methods and challenges. Group Decision and Negotiation **10**(2), 199–215 (2001)
62. Kirschner, P.A., Shum, S.J.B., Carr, C.S. (eds.): Visualizing Argumentation: Software Tools for Collaborative and Educational Sense-Making. Springer, London, UK (2003). URL http://oro.open.ac.uk/12107/
63. Langley, P., Simon, H.A.: Applications of machine learning and rule induction. Commun. ACM **38**(11), 54–64 (1995). DOI http://doi.acm.org/10.1145/219717.219768
64. Leite, J., Martins, J.: Social abstract argumentation. In: T. Walsh (ed.) IJCAI 2011, Proceedings of the 22nd International Joint Conference on Artificial Intelligence, Barcelona, Catalonia, Spain, July 16-22, 2011, pp. 2287–2292. IJCAI/AAAI (2011)
65. Lucero, M.G., Chesñevar, C., Simari, G.R.: On the accrual of arguments in defeasible logic programming. In: C. Boutilier (ed.) IJCAI, pp. 804–809 (2009)
66. Mackenzie, J.: Four dialogue systems. Studia Logica **49**(4), 567–583 (1990)
67. Matt, P.A., Morge, M., Toni, F.: Combining statistics and arguments to compute trust. In: W. van der Hoek, G.A. Kaminka (eds.) Proceedings of the 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010) (2010)
68. Matt, P.A., Toni, F.: A game-theoretic measure of argument strength for abstract argumentation. In: 11th European Conference on Logics in Artificial Intelligence (2008)
69. Matt, P.A., Toni, F., Stournaras, T., Dimitrelos, D.: Argumentation-based agents for eprocurement. In: M. Berger, B. Burg, S. Nishiyama (eds.) Proceedings of the 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)- Industry and Applications Track, pp. 71–74 (2008)
70. Matt, P.A., Toni, F., Vaccari, J.: Dominant decisions by argumentation agents, argumentation in multi-agent systems. In: P. McBurney, S. Parson, I. Rawan, N. Maudet (eds.) ArgMAS 2009 (2009)
71. McGinnis, J., Stathis, K., Toni, F.: A formal model of agent-oriented virtual organisations and their formation. Multiagent and Grid Systems **7**(6), 291–310 (2011)
72. Mercier, H., sperber, D.: Why do humans reason? arguments for an argumentative theory. Behavioral and Brain Sciences **34**(2), 57–747 (2011)
73. Mitchell, T.: Machine Learning. McGraw-Hill Education (ISE Editions) (1997). URL http://www.amazon.ca/exec/obidos/redirect?tag=citeulike09-20\&path=ASIN/0071154671
74. Modgil, S.: Reasoning about preferences in argumentation frameworks. Artificial Intelligence **173**(9-10), 901–934 (2009)
75. Modgil, S., Bench-Capon, T.: Metalevel argumentation. Journal of Logic and Computation pp. 959–1003 (2011)
76. Modgil, S., Caminada, M.: Proof theories and algorithms for abstract argumentation frameworks. In: I. Rahwan, G. Simari (eds.) Argumentation in AI, pp. 105–129. Springer-Verlag (2009)
77. Modgil, S., Prakken, H.: Revisiting preferences and argumentation. In: 22nd International Joint Conference on Artificial Intelligence, pp. 1021–1026 (2011)
78. Moguillansky, M., Rotstein, N., Falappa, M., García, A., Simari, G.R.: Argument Theory Change: Revision Upon Warrant. In: Proceedings of The Twenty-Third Conference on Artificial Intelligence, AAAI 2008, pp. 132–137 (2008)

79. Moguillansky, M., Rotstein, N., Falappa, M., García, A., Simari, G.R.: Argument theory change through defeater activation. In: P. Baroni, F. Cerutti, M. Giacomin, G.R. Simari (eds.) COMMA, *Frontiers in Artificial Intelligence and Applications*, vol. 216, pp. 359–366. IOS Press (2010)
80. Možina, M.: Argument based machine learning. Ph.D. thesis, University of Ljubljana: Faculty of Computer and Information Science, Ljubljana, Slovenia (2009)
81. Možina, M., Guid, M., Krivec, J., Sadikov, A., Bratko, I.: Fighting knowledge acquisition bottleneck with argument based machine learning. In: The 18th European Conference on Artificial Intelligence (ECAI), pp. 234–238. Patras, Greece (2008)
82. Možina, M., Guid, M., Krivec, J., Sadikov, A., Bratko, I.: Learning to explain with ABML. In: Proc. of the 5th International Workshop on Explanation-aware Computing (ExaCt'2010), ISNN 1613-0073, pp. 37–49. CEUR-WS.org (2010)
83. Možina, M., Žabkar, J., Bratko, I.: Argument based machine learning. Artificial Intelligence **171**(10/15), 922–937 (2007)
84. Niemelä, I.: Logic programming with stable model semantics as a constraint programming paradigm. Ann. Math. Artif. Intell. **25**(3–4), 241–273 (1999)
85. Pang, B., Lee, L.: Opinion mining and sentiment analysis. Foundations and Trends in Information Retrieval **2**(1-2), 1–135 (2008)
86. Podlaszewski, M., Caminada, M., Pigozzi, G.: An implementation of basic argumentation components (demonstration). In: Proceedings AAMAS 2011, pp. 1307–1308 (2011)
87. Prakken, H.: Modelling reasoning about evidence in legal procedure. In: Proceedings of the Eighth International Conference on Artificial Intelligence and Law (ICAIL-01), pp. 119–128. ACM Press, New York (2001)
88. Prakken, H.: Coherence and flexibility in dialogue games for argumentation. Journal of Logic and Computation **15**(6) (2005)
89. Prakken, H.: An abstract framework for argumentation with structured arguments. Argument and Computation **1**(2), 93–124 (2010)
90. R, R.G., Hitchcock, D., McBurney, P., Verheij, B.: Argumentation Machines. New Frontiers in Argument and Computation, chap. Decision Support for Practical Reasoning: a theoretical and computational perspective., pp. 55–84. Kluwer Academic Publishers, Dordrecht (2003)
91. Rahwan, I., Ramchurn, S.D., Jennings, N.R., McBurney, P., Parsons, S., Sonenberg, L.: Argumentation-based negotiation. The Knowledge Engineering Review **18**(4), 343–375 (2004)
92. Rahwan, I., Reed, C.: The argument interchange format. In: I. Rahwan, G. Simari (eds.) Argumentation in Artificial Intelligence, pp. 383–402. Springer (2009)
93. Rahwan, I., Zablith, F., Reed, C.: Laying the foundations for a world wide argument web. Artificial Intelligence **171**, 897–921 (2007)
94. Reed, C., Rowe, G.: Araucaria: Software for Argument Analysis, Diagramming and Representation. International Journal of AI Tools **13**(4), 961–980 (2004)
95. Reed, C., Wells, S., Budzynska, K., Devereux, J.: Building arguments with argumentation : the role of illocutionary force in computational models of argument. In: Proceedings of COMMA 2010 (2010)
96. Rotstein, N., Moguillansky, M., Falappa, M., García, A., Simari, G.R.: Argument Theory Change: Revision upon Warrant. In: Proceedings of The Computational Models of Argument, COMMA 2008, pp. 336–347 (2008)
97. Sadikov, A., Možina, M., Guid, M., Krivec, J., Bratko, I.: Automated chess tutor. In: Proceedings of the 5th International Conference on Computers and Games (2006)
98. Schneider, J., Passant, A., Groza, T., Breslin, J.G.: Argumentation 3.0: how semantic web technologies can improve argumentation modeling in web 2.0 environments. In: P. Baroni, F. Cerutti, M. Giacomin, G.R. Simari (eds.) Computational Models of Argument: Proceedings of COMMA 2010, Desenzano del Garda, Italy, September 8-10, 2010, *Frontiers in Artificial Intelligence and Applications*, vol. 216, pp. 439–446. IOS Press (2010)
99. Shafer, G.: Probability judgment in artificial intelligence. In: L.N. Kanal, J.F. Lemmer (eds.) Proceedings of the First Annual Conference on Uncertainty in Artificial Intelligence (UAI'85), pp. 127–136. Elsevier (1985)

100. Snaith, M., Lawrence, J., Reed, C.: Mixed initiative argument in public deliberation. In: D.C. F., M. A., P. C. (eds.) From e-Participation to Online Deliberation, Proceedings of OD2010 (2010)
101. Tolchinsky, P., Cortés, U., Modgil, S., Caballero, F., Lopez-Navidad, A.: Increasing the availability of human organs for transplantation through argumentation based deliberation among agents. IEEE Special Issue on Intelligent Agents in Healthcare **21(6)**, 30–37 (2006)
102. Tolchinsky, P., Modgil, S., Atkinson, K., McBurney, P., Cortes, U.: Deliberation dialogues for reasoning about safety critical actions. Journal of Autonomous Agents and Multi-Agent Systems (JAAMAS) **25**, 209–259 (2012)
103. Tolchinsky, P., Modgil, S., Cortés, U., Sánchez-Marré, M.: Cbr and argument schemes for collaborative decision making. In: Proc. 1st International Conference on Computational Models of Argument, pp. 71–82. IOS Press, Liverpool, UK (September, 2006)
104. Toni, F.: Argumentative KGP agents for service composition. In: M. Balduccini, C. Baral (eds.) Proc. AITA08, Architectures for Intelligent Theory-Based Agents, AAAI Spring Symposium. Stanford University (2008)
105. Toni, F.: Argumentative agents. In: Proceedings of the International Multiconference on Computer Science and Information Technology, pp. 223–229. IEEE (2010)
106. Toni, F., Grammatikou, M., Kafetzoglou, S., Lymberopoulos, L., Papavassileiou, S., Gaertner, D., Morge, M., Bromuri, S., McGinnis, J., Stathis, K., Curcin, V., Ghanem, M., Guo, L.: The ArguGRID platform: An overview. In: J. Altmann, D. Neumann, T. Fahringer (eds.) Proceedings of Grid Economics and Business Models, 5th International Workshop (GECON 2008), *Lecture Notes in Computer Science*, vol. 5206, pp. 217–225. Springer (2008)
107. Toni, F., Sergot, M.: Argumentation and ASP. In: LP, KR, and NMR: Essays in Honor of Michael Gelfond. Springer (2011)
108. Toni, F., Torroni, P.: Bottom-up argumentation. In: S. Modgil, N. Oren, F. Toni (eds.) First International Workshop on Theory and Application, TAFA 2011, Barcelona, Spain, July 16-17, 2011, Revised Selected Papers, *Lecture Notes in Computer Science*, vol. 7132, pp. 249–262. Springer (2012)
109. Torroni, P., Gavanelli, M., Chesani, F.: Arguing on the semantic grid. In: I. Rahwan, G. Simari (eds.) Argumentation in AI, pp. 423–441. Springer (2009). DOI 10.1007/978-0-387-98197-0_21
110. Torroni, P., Prandini, M., Ramilli, M., Leite, J., Martins, J.: Arguments against the troll. In: Proceedings of the Eleventh AI*IA Symposium on Artificial Intelligence, Brescia, December, pp. 232–235. Arti Grafiche Apollonio, Brescia (2010)
111. Toulmin, S.: The Uses of Argument. Cambridge University Press (1958)
112. van Veenen, J., Prakken, H.: A protocol for arguing about rejections in negotiation. In: S. Parsons, N. Maudet, P. Moraitis, I. Rahwan (eds.) Proceedings of the 2nd International Workshop on Argumentation in Multi-Agent Systems (ArgMAS 2005), affiliated to AAMAS 2005, *Lecture Notes in Computer Science*, vol. 4049, pp. 138–153. Springer (2006)
113. Verheij, B.: Artificial argument assistants for defeasible argumentation. Artif. Intell. **150**(1-2), 291–324 (2003)
114. Verheij, B.: A labeling approach to the computation of credulous acceptance in argumentation. In: M.M. Veloso (ed.) Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI 2007), pp. 623–628 (2007)
115. Walton, D.N.: Argument Schemes for Presumptive Reasoning. Lawrence Erlbaum Associates, Mahwah, NJ, USA (1996)
116. Walton, D.N., Krabbe, E.C.W.: Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning. SUNY Series in Logic and Language. State University of New York Press, Albany, NY, USA (1995)
117. Walton, D.N., Reed, C., Macagno, F.: Argumentation Schemes. Cambridge University Press (2008)
118. Webb, G.I., Wells, J., Zheng, Z.: An experimental evaluation of integrating machine learning with knowledge acquisition. Mach. Learn. **35**(1), 5–23 (1999). DOI http://dx.doi.org/10.1023/A:1007504102006

119. Wells, S., Gourlay, C., Reed, C.: Argument blogging. In: 9th International Workshop on Computational Models of Natural Argument (2009)
120. Wooldridge, M.: Properties and complexity of some formal inter-agent dialogues. Journal of Logic and Computation **13**, 347–376 (2003)
121. Yu, B., Singh, M.P.: Distributed reputation management for electronic commerce. Computational Intelligence **18**(4), 535–549 (2002)