Full Acceptance via Argumentation

Daniela Carbogim[1] and Renata Wassermann[2]

[1] Division of Informatics, University of Edinburgh
danielac@dai.ed.ac.uk
[2] Department of Computer Science, University of São Paulo
renata@ime.usp.br

Abstract. Any rational agent must have a strategy for deciding whether a piece of uncertain information is acceptable or not. In this paper we argue for the use of argumentation theory to solve this problem. The decision process is seen as an internal argumentation, whereby the agent weights evidence for and against some piece of information. We work within a framework for belief revision which takes into account the agent's limited resources, and apply a formal framework for argumentation to the problem of fully accepting information.
Keywords: argumentation, belief change, bounded rationality.

1 Introduction

The problem of belief revision, i.e. how the beliefs of an agent should change in the presence of new information, has been addressed by various authors. In most approaches, especially those following the AGM paradigm [1], the agents are idealised in that they are assumed to have perfect recall and to hold only consistent beliefs, which are furthermore assumed to be closed under logic consequence.

In [13], Wassermann has developed a model for belief revision which takes into account the effects of both limited memory and limited capacities of inference. The belief state of an agent is represented by a structure that distinguishes different kinds of beliefs: beliefs that are explicit or basic, beliefs that are implicit or merely derived, and beliefs that are active or in use. Besides the beliefs of the agent, the structure also represents *provisional beliefs*, i.e. sentences for which the agent has some evidence but is not yet sure whether to accept them or not. In [6] and [14], some ways of deciding which beliefs were active during an operation of belief change were explored. In this paper we turn to another question left open by the model, namely how to decide if a provisional belief should be accepted.

Incoming information is usually given the highest priority, so that if a contradiction arises, some of the previous beliefs have to be given up. In approaches to non-prioritised belief revision [5], where the new piece of information does not have the highest priority, the decision of whether to accept or not new information is taken by extra-logical means such as selection functions or incision functions, but there is no real recipe on how to choose these functions.

A different idea was explored in [12] – that of using argumentation theory for deciding if new information is acceptable. According to Dung [4], reasoning about one's own beliefs is like performing an internal argument. A sentence is believable "if it can be argued successfully against attacking arguments" [4]. Our concept of provisional beliefs is based on Harman's idea of *tentative hypotheses.*

In order to be fully accepted, a tentative hypothesis has to survive the best attempts to refute it [7]. In our case, "best attempts" are as good as the agent is capable given his limitations.

Argumentation then provides a generic account for a style of reasoning. The idea behind it is that we can make and justify decisions by constructing and considering arguments relevant to conflicting alternatives. As a reasoning technique, argumentation is less concerned with restoring consistency. Instead it suggests that conflict and inconsistency can offer insights into rational processes and therefore should not be necessarily eradicated.

Here we elaborate on the initial ideas in [12] and specify an argument-based framework for deciding if a provisional belief should be accepted in the resource-bounded model in [13]. Although a number of theoretical frameworks for argumentation exist, most are not mature enough to allow specific applications to be instantiated and developed in a systematic way [2]. To provide a practical, yet non *ad hoc* framework, we need to focus on the properties and characteristics of our belief revision model, and on how these will affect our design choices.

This paper is organised as follows. Section 2 outlines the belief revision model in [13], and the problem we are tackling. In Section 3 we discuss argumentation as technique for reasoning, and in Section 4 we describe an argument-based approach to our problem. Section 5 summarises our contributions.

2 A Resource-Bounded Framework for Belief Revision

In this model we distinguish between different types of beliefs. Following Harman [7], we assume that there are beliefs that are explicitly represented. We call *implicit beliefs* those beliefs that can be inferred from the set of the agent's *explicit beliefs*, according to the agent's logical ability. We will not concentrate in one particular inference operator, but use *Inf* to denote what an agent can infer in one step. The set of implicit beliefs is given by what the agent would be able to infer if he was given unlimited time, i.e. the result of applying *Inf* an unlimited number of times.

Not all beliefs are available at the same time. We call *active beliefs* the set containing beliefs that are available for use and things about which the agent is not yet sure. These last are called *provisional beliefs*. Provisional beliefs are outside the set of explicit beliefs, and are still under inspection.

A belief state is a structure $\beta = \langle E, \mathit{Inf}, A \rangle$, where E is the set of the agent's explicit beliefs, *Inf* is the agents inference function and A is the set of the agent's active beliefs. The set of implicit beliefs is given by: $I = \mathit{Inf}^*(E) = \mathit{Inf}(E) \cup \mathit{Inf}^2(E) \cup \ldots$.. Given our representation of belief state, we now define operations that can be applied to belief states to modify them.

The set of active beliefs is based on the idea of a short-term memory. Active beliefs can be forgotten or stored as explicit (but inactive) beliefs. Since the set of active beliefs is assumed to be very limited in size[1], there must be a mechanism that, in cases of overflow, selects which beliefs will be forgotten or stored.

The first operation we define is similar to AGM expansion in the sense that it consists in simply adding new information to a set without checking for con-

[1] By the size of a set of formulae we mean something like its complexity, e.g. the number of occurrences of atoms. The sets $\{p, q\}$ and $\{p \wedge q\}$ have the same size.

sistency. But the operation takes the limited size of the set into account. When trying to add something to a set that is already at its maximum size, some elements of the set have to be given up. This can be seen as a kind of "forgetting".

If X is a set with maximum size k and α is an element we want to add to X, then $X \cup^* \{\alpha\} = X' \cup \{\alpha\}$, where $X' \subseteq X$, $|X'| < k$. Note that this operation reduces to a simple union as long as the set is not "full". Since the size k of the set is given as a parameter, the operation is more accurately denoted as $\cup^*_k$. When the set is already at its maximum size, something has to be discarded. If the set X is ordered (for example by the last time the beliefs were recalled), we can stipulate that the minimal elements of the set are the first to be dismissed, i.e. we want to ensure that if an element is dismissed, then there is no other element which is retained and that is less than the dismissed one in the order: $\forall y(y \in X \setminus X' \rightarrow \neg\exists x(x \in X' \wedge x < y))$.

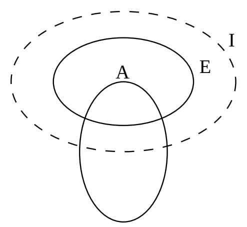

Fig. 1. *Structure of an agent's beliefs.*

Definition 1 *Let $\langle E, Inf, A\rangle$ be a belief state and α a formula. We define the following operations on $\langle E, Inf, A\rangle$ (we will omit the second argument Inf since the operations defined do not affect it):*

1. *Observation ($\circ_o$): adds an external input to the set of active beliefs.*
   $\langle E, A\rangle \circ_o \alpha = \langle E, A \cup^* \{\alpha\}\rangle$
2. *Retrieval ($\circ_r$): retrieves an explicit belief into the set of active beliefs.*
   $\langle E, A\rangle \circ_r \alpha = \begin{cases} \langle E, A \cup^* \{\alpha\}\rangle, & \text{if } \alpha \in E \\ \langle E, A\rangle & \text{otherwise} \end{cases}$
3. *Acceptance ($\circ_a$): makes an active belief explicit.*
   $\langle E, A\rangle \circ_a \alpha = \begin{cases} \langle E \cup^* \{\alpha\}, A \setminus \{\alpha\}\rangle, & \text{if } \alpha \in A \\ \langle E, A\rangle & \text{otherwise} \end{cases}$
4. *Inference ($\circ_i$): infers something from active beliefs.*
   $\langle E, A\rangle \circ_i \alpha = \begin{cases} \langle E, A \cup^* \{\alpha\}\rangle, & \text{if } \alpha \in Inf(A) \\ \langle E, A\rangle & \text{otherwise} \end{cases}$
5. *Doubting ($\circ_d$): a belief that was accepted is questioned, becoming provisional.*
   $\langle E, A\rangle \circ_d \alpha = \begin{cases} \langle E \setminus \{\alpha\}, A\rangle, & \text{if } \alpha \in A \cap E \\ \langle E, A\rangle & \text{otherwise} \end{cases}$
6. *Rejection ($\circ_c$): rejects an active belief.*
   $\langle E, A\rangle \circ_c \alpha = \begin{cases} \langle E, A \setminus \{\alpha\}\rangle, & \text{if } \alpha \in A \\ \langle E, A\rangle & \text{otherwise} \end{cases}$

The six operations defined above can be combined to model more complex ones. As an example of such a composition, consider what happens when an agent gets new information via observation. The belief first comes into the set of active beliefs through the operation $\circ_o$ and then the agent may accept it ($\circ_a$).

2.1 The Problem: Accepting Provisional Beliefs

At this point we turn to the problem of deciding whether a provisional belief should be accepted, i.e. made explicit. Given a belief state $\langle E, A\rangle$, the set of provisional beliefs is a subset of $A \setminus I$. In terms of the operations in Definition 1, there are two ways in which a sentence α can become a provisional belief:

– α is acquired by an observation ($\circ_o$);
– α was previously accepted (explicit) and is questioned by doubting ($\circ_d$).

The question we address in this paper can be stated as follows. *Let $\langle E, A\rangle$ is a belief state, and $\alpha \in A \setminus I$ a provisional belief upon which the agent wants to decide. How to determine whether to accept or reject α (i.e. whether to perform an acceptance operation $\langle E, A\rangle \circ_a \alpha$, or a rejection operation $\langle E, A\rangle \circ_c \alpha$)?*

We now take a look at argumentation as an answer to our question. We first take a bird's eye view on formal theories of argumentation and applications. Next we give careful consideration to the properties in our model—such as resource-boundedness, belief structure and change operations—in order to specify an appropriate argument-based framework for deciding upon a provisional belief.

3 Formal Models of Argumentation

As discussed in the Introduction, argumentation gives a generic architecture for a particular style of reasoning, which consists in constructing and comparing arguments in order to make decisions. Formal theories of argumentation make use of formal languages and inference techniques to capture such processes of argument exchange, where arguments correspond to proofs generated from an underlying set of formulae via an inference relation. What is interesting about argumentation is that it explores aspects of practical reasoning that are not always addressed by conventional reasoning theories. For example, instead of focusing on restoring consistency, argumentation is based on the notion of acceptability: a claim is acceptable on the basis of the arguments that are relevant to it.

This general paradigm of argument-based reasoning has proved applicable to a variety of tasks involving inconsistent and incomplete information. In fact, in a recent survey [2] we have identified the sorts of generic problems that have been solved by argumentation methodologies, which range from non-monotonic and defeasible reasoning to design. Because of the intrinsic non-monotonic nature of our problem, we pay particular attention to the use of argumentation in non-monotonic and defeasible theories, i.e. to *defeasible argumentation*[2].

Defeasible argumentation systems are usually intended to determine if an argument is acceptable based on its relations to other arguments (generated from a static knowledge base). The status of an argument depends on the set of all arguments, and can be specified in two ways: *declaratively*, by defining a class of acceptable arguments; and *procedurally*, via a procedure for determining if an argument is in this class. We initially thought of applying one such system to our problem, but the assumption that *all* arguments need to be considered clashes with the resource-bounded, dynamic nature of the belief model in Section 2.

A different view of procedural models was summarised by Loui in [8], where he argues that what makes beliefs rational is not only their relations to other beliefs, but also the way in which they are built as the outcome of deliberative processes. Loui gives an account of defeasible argumentation as resource-bounded, dialectic disputation protocols based on notions such as which parties are involved; what are the possible moves for each party; how moves affect the

[2] For a comprehensive account on defeasible argumentation systems see [11, 3].

outcome; how to determine if a disputation has finished; and if it has been won or lost. For the outcome to be rational, such protocols must be *fair* (e.g. parties get the same amount of resources) and *effective* (e.g. when a conclusion is established, it means that maximum resources were used in unsuccessful criticisms).

Loui [8] defines a very general framework for argumentation that has to be *filled in* in order to model particular kinds of disputation, and presented some instances of protocols based on an explicit defeasible reason relation and on the specificity principle for deciding between conflicting beliefs. This framework has been considered in [12] as a potential solution for the problem of full acceptance. However, *filling it in* in an adequate way is not always a straightforward task. As argued in [2], "since argumentation, in automated forms, is relatively new there does not yet exist methods for guiding application of architectures to problems, and the focus has been on more abstract argumentation theory".

In this paper we focus on the properties of the belief model in Section 2 so as to present an argumentation framework that is suitable to our problem. To guide this task we use a generic conceptual model that underlies most defeasible argumentation systems [9], and which we discuss below.

3.1 A Conceptual Framework for Defeasible Argumentation

Prakken [9] has identified five main concepts behind defeasible argumentation formalisms. We briefly discuss these concepts (based on [11]), before presenting a general account of disputation protocols which will be the basis of our proposal.

Underlying Logic The underlying logic is essentially the logical system defining a *monotonic* consequence relation as the basis for deriving arguments.

Arguments Arguments correspond to proofs in the underlying formal system.

Conflict Argumentation implies disagreement, captured by the idea of conflict.

Defeat Since the underlying logic is monotonic, the addition of new information does not invalidate existing arguments, so conflicting arguments may coexist. The non-monotonic character of argumentation arises because some arguments may be preferred over others. Such preference criteria may not always exist, but if they do, these are likely to be domain specific.

Acceptability Status Intuitively, an argument that defeats a conflicting argument but is also defeated by a third one is not acceptable. An argument acceptability status is often determined by processes that, instead of looking at conflicting arguments alone, take into account all the other relevant arguments.

In brief, the main purpose of an argument is to give a justification. But arguments can also be used to attack others, for instance in the form of counter arguments supporting conflicting views. This role is concerned not with individual justifications, but with the relationship between these, explored by means of processes of argument exchange, i.e. argumentation protocols. Based on [8, 10], we now give an abstract account of dialectical protocols that are intended to determine the acceptability status of an argument Arg.

Definition 2 *Let $\mathcal{P}$ be the set of participants in a dialectical system, namely a proponent P and an opponent O. A dialectical argument about Arg is a sequence:*

$$\mathcal{D} = \langle (P_1, Arg_1), (P_2, Arg_2), ..., (P_n, Arg_n) \rangle,$$

specified by the following functions defined over the subsequence of moves advanced so far at each step $k \geq 1$*, namely* $\mathcal{D}_k = \langle (P_1, Arg_1), ..., (P_{k-1}, Arg_{k-1}) \rangle$*:*

- $P_k = PartyToMove(\mathcal{D}_k) \in \mathcal{P}$*, and* $PartyToMove(\emptyset) = P$*;*
- $Arg_k \in M_k$*,* $M_k = LegalMoves(P_k, \mathcal{D}_k)$ *is the set of arguments that* P_k *can advance to respond to the last move* Arg_{k-1}*, and* $LegalMoves(P, \emptyset) = \{Arg\}$

If $M_{n+1} = \emptyset$*, then* P_n *is considered the "winning" party. That is, if* $P_n = P$ *then the proponent has won the disputation and* Arg *is said to be acceptable.*

Important aspects in this definition remain unspecified, especially how to characterise the set of possible argument moves and how to select which argument (from this set) to advance. In accounts such as [8], where participants have limited shares of resources at their disposal and maybe access to distinct bodies of information, the argument selection mechanism is bounded by the available resources and information. That means that the function *LegalMoves* can take into account limited memory and limited capacities of inference, hence the set M_{n+1} may be empty not only because the party ran out of arguments but also because it ran out of resources, or time. This is a much more realistic situation, but for the outcome of such processes to be adequate this mechanism must be carefully designed. In Loui's words, protocols must be *fair* and *effective*.

More recently, Prakken [10] has also focused on the study of dialectical protocols, but from a slightly different perspective than Loui's. Rather than considering partial computation and limited resources, Prakken is interested in cases where new information is added during the process, and in characterising the properties that make protocols appropriate in these situations (e.g. if a participant could have advanced an attack, it had the chance to do so during the argumentation). In Prakken's words, protocols must be *fair* and *sound*.

Both resource-boundedness and changes to the underlying base are relevant to our problem, but we need to determine exactly *how*. So we turn again to our belief model so as to appropriately instantiate this argumentation framework.

4 An Argumentation Framework for Accepting Beliefs

In [12], Wassermann has characterised and illustrated the type of behaviour expected from an argumentation framework for accepting provisional beliefs, and proposed the use of Loui's framework to describe this behaviour. Our aim in this work is to follow up on that proposal and fully describe an argument-based technique that captures that intended behaviour.

In our case, argumentation is an internal process where a single agent plays the role of the proponent and the opponent, analysing the arguments for and against a given provisional belief. Since we are dealing with resource-bounded agents, this internal argumentation will not always succeed in examining all reasons for accepting or rejecting a belief. By defining a protocol for this process, we have to take care that the outcome can be considered fair. We start by instantiating the components in the conceptual framework of Section 3.1.

Underlying Logic The underlying logical system is defined by the language of propositional calculus, and inference rules *modus ponens* and *modus tollens*. Given a belief state $\langle E, A \rangle$, E and A are sets of formulae in this system.

Argument Formal arguments correspond to inference processes. In this model, an inference process is a sequence of change operations for activating beliefs, so it essentially involves changes to the belief states. Given a state $\langle E_0, A_0\rangle$, an argument for a sentence α is characterised as a minimal sequence

$$\langle E_0, A_0\rangle, ..., \langle E_m, A_m\rangle,$$

such that for each $0 < j \leq m$, $\langle E_j, A_j\rangle = \langle E_{j-1}, A_{j-1}\rangle *_j \alpha_j$, where $*_j$ is an operation and $\alpha = \alpha_m \in A_m$.

We can make some simplifying assumptions by restricting the types of operations that can be applied for constructing an argument. First, because we are applying argumentation to decide whether to *accept* ($\circ_a$) or *reject* ($\circ_c$) a sentence, we can assume that these operations are not part of an argument. Second, the *doubting* ($\circ_d$) operation is relevant to argumentation (since a belief can be questioned if an argument against it is presented), but it does not seem to be part of a demonstrative argument itself. Third, observations ($\circ_o$) are performed by external sources and, although they bear on the argumentation as they introduce new provisional beliefs, they are independent from the process. In this way we restrict the types of operations to *retrieval* and *inference*. Consequently, the set of explicit beliefs does not change during the construction of an argument, which is now represented by

$$\langle E, A_0\rangle, ..., \langle E, A_m\rangle,$$

where each $*_j$ is either $\circ_r$, or $\circ_i$. This argument is represented by an (ordered) set Arg_α of the beliefs used in the inference of α. Note that inference operations may be based on beliefs that were already in A_0 (and which remained active until inference was applied). Let $A'_0 \subseteq A_0$ be the set of these beliefs. Then elements in Arg_α are drawn from $A'_0 \cup \{\alpha_j | 0 < j \leq m\}$, and $Arg_\alpha \subseteq^* A_m$[3].

Conflict An argument Arg_α is said to *attack* an argument Arg_β if they are in direct conflict, i.e. if $\alpha = \neg\beta$.

Defeat In our model we do not assume any ordering on the set of active beliefs, apart from that of *recency*[4]. In the context of conflict, every argument that attacks another is "strong enough" to defeat it.

Note that an argument for $\neg\alpha$ serves as evidence against α, so the agent can question α (perform a doubting operation) as a consequence of this defeat.

Acceptability Status A sentence can have two possible status – *acceptable* or *non-acceptable* – which is determined by the protocol defined below.

We now instantiate the generic protocol in Definition 2. We can assume that there is a recency ordering in which provisional beliefs will be tested via argumentation. Note also that this process will involve change operations to the initial belief state $\langle E, A\rangle$ into some other state $\langle E', A'\rangle$, so a provisional belief α will be accepted or rejected with respect to this latter state rather than to the initial one. Moreover, α must remain active during the whole argumentation, i.e. it cannot be dismissed due to overflow in the set of active beliefs. Hence during this process α keeps a fixed highest priority in the recency ordering in A.

[3] The relation $X \subseteq^* Y$ is based on the usual subset relation restricted to the size of Y. If $|X| \leq |Y|$, then $X \subseteq^* Y = X \subseteq Y$, otherwise, $X \subseteq^* Y = X' \subseteq Y$, where $X' \subseteq X$ and $|X'| \leq |Y|$.

[4] Beliefs that are used first are the first to be forgotten in case of overflow. However, if an active belief is reused, it becomes more recent and changes place in the order.

Definition 3 *Let $\langle E, A\rangle$ be a belief state and $\alpha \in A$ be a provisional belief upon which the agent wants to decide. An argumentation about α is a sequence*

$$\langle E, A\rangle = \langle E_1, A_1\rangle, \langle E_2, A_2\rangle, ..., \langle E_r, A_r\rangle = \langle E', A'\rangle,$$

of changes to the initial belief state represented by a process of argument exchange

$$\mathcal{D} = \langle (P_1, Arg_1), (P_2, Arg_2), ..., (P_n, Arg_n)\rangle, \text{ where}$$

- $P_1 = P$;
- *if* $P_k = P$, $k \geq 1$, *then* $P_{k+1} = O$ *and vice versa (the function PartyToMove is about alternating turns between the two roles played by the agent);*
- $Arg_1 = \{\alpha\} \subseteq A_1$, *and also* $\alpha \in A_k$ *for* $1 \leq k \leq r$;
- $Arg_k \in M_k$, $M_k = LegalMoves(P_k, \mathcal{D}_k)$ *is the set of arguments that* P_k *can advance to respond to the last move* Arg_{k-1}.

If $M_{n+1} = \emptyset$, *then* P_n *is considered the "winning" party. If the proponent wins* ($P_n = P$) *then the agent accepts* α *and rejects* $\neg\alpha$. *If the opponent wins* ($P_n = O$) *then the agent rejects* α, *but it does not accept* $\neg\alpha$ *(it can be ignorant about* α*).*

The crucial step is to define the function *LegalMoves*, which must take into consideration the current set of active beliefs and some limitation on the number of basic operations that an agent can perform during a disputation. In our belief model these are clearly limited resources. Also, since an agent is playing the roles of both proponent and opponent, we can assume that the set of active beliefs is a shared resource that both can access as a whole. In this context, *LegalMoves* represent the set of possible arguments that attack (and defeat) any argument advanced by the other party and that is still relevant to the process.

But what exactly does it mean for an argument to be still relevant? As the argumentation proceeds, each participant activates a number of beliefs, delineating a subset of active beliefs that were used to support its views and can therefore be attacked. Since this subset is bounded by the size of the set of active beliefs, it can increase as new beliefs are activated or decrease as other beliefs are forgotten. In fact, let $\langle E_j, A_j\rangle$, $1 \leq j \leq r$ be the corresponding belief state at step k in the argumentation, just after P_k has advanced Arg_k. At this point, the set of sentences that have been asserted by P_k as it generated other arguments can be characterised by the set $X_k = \bigcup_{i \leq k, P_i = P_k} Arg_i$. Hence, the set of arguments from P_k that can be attacked at the next step $k+1$ is defined by $\mathcal{A}_k^{P_k} = X_k \cap A_j$. An interesting point is that we do not take into consideration the structure of arguments, but just the sentences that were derived from them. This is because the agent incorporates all the beliefs it infers, and this is something characteristic from our belief model which we have to account in the argumentation framework.

Intuitively, the set M_k is then composed of all arguments that can be generated so that they attack (and defeat) some argument from $\mathcal{A}_k^{P_k}$. In our case, however, generating an argument involves changes in the set of active beliefs, so constructing the whole set M_k and then selecting an attack from it does not seem to be the most intuitive strategy. What we do is to construct one such argument, and then advance it. If P_k cannot derive an argument, either because M_k is empty or because its resources are exhausted, then it looses.

Note that we can apply certain strategies to the construction of an argument, especially for selecting which sentence to attack and which beliefs to retrieve. Intuitively, because arguments are generated from A, they are based both on beliefs that are accepted (explicit) and on beliefs that are provisional. An argument that is based on an accepted belief is "stronger" than arguments that are based only on provisional ones. A suitable strategy would be attempting to use as many accepted beliefs as possible when constructing an argument (as they are stronger), but selecting provisional beliefs to attack (as they are weaker).

Trying to construct an argument may not be always be an easy (and successful) task. In fact, it can lead to a lot of resource waste, i.e. to the application of irrelevant operations. And as a participant spends resources, the set of active beliefs is changing. For instance, as a party unsuccessfully tries to construct an argument, the subset representing the sentences it can attack may be completely forgotten, so eventually the party looses because it could not generate a suitable attack and now there are no more choices.

Note that the relation between the sequence of belief states and the sequence of arguments in a disputation is not *one-to-one*. We can identify subsequences of belief states that represent "pure" arguments, but other types of operations can also be applied during the argumentation. Some of these changes refer to doubting (after a sentence has been successfully attacked), others are performed as the agent unsuccessfully attempts to generate an argument.

All these issues lead to the following question: *can we precisely describe what are the properties that this protocol should have for it to be considered fair?* It is clear that we cannot apply the types of characterisation given by Loui [8] and Prakken [10], as the nature of our problem is different from both. The first idea is to divide the resources per player, so that one cannot erase the other's retrieved beliefs. However, it comes down to the size of the resources as well as to the size of A. Even though E has also limited size, no belief is added to it during the argumentation. In any case, the appropriate size of A and the amount of resources can only be defined empirically, and it is likely that it will be specific to the domain of application.

We are currently designing an experimental system in which we can perform empirical tests based on a number of parameters, such as the size of A, the amount of resources, strategies for selecting a sentence to attack and strategies for generating arguments. We hope to form hypotheses about the relation between these parameters and the appropriateness of the protocol, thus strengthening our understanding of real argumentation applications. Note that this idea is actually in line with Dung's remark that "correctness" of an argumentation theory for capturing the idea of acceptability cannot be "proved" formally; rather we can prove "appropriateness" by giving relevant and convincing examples [4].

5 Conclusion

In this paper we have applied a formal argument model to the problem of accepting information. We first presented a general argumentation framework based on Prakken's [10] and Loui's [8] models, and then we instantiated it to our particular case. In doing so, we could describe the argumentation process by means of the basic operations for resource-bounded belief revision proposed in [13].

To summarise, the main contribution of the present work is a new application of an existing technique, namely argumentation. Because existing argumentation formalisms are based on assumptions that do not fit the characteristics of our problem, we had to define an argument-based approach that would take these into account. So another contribution of this work was the systematic instantiation of this approach from a very abstract account of argumentation. This instantiation may serve as guidance for future application of argumentation theories to concrete problems.

Further work includes a series of experiments and empirical tests in order to form hypotheses about appropriate values for the parameters in our framework. In particular, we would like to examine adequate strategies of argument selection and to search for a good size for the set A of active beliefs.

References

1. C. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change. *Journal of Symbolic Logic*, 50(2):510–530, 1985.
2. D. Carbogim, D. Robertson, and J. Lee. Argument-based applications to knowledge engineering. *The Knowledge Engineering Review*, forthcoming, 2000.
3. C. I. Chesñevar, A. G. Maguitman, and R. P. Loui. Logical models of argument. *ACM Computing Surveys*, `www.cs.wustl.edu/~loui/survey.ps`, 1999.
4. P. M. Dung. 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.
5. S. O. Hansson, editor. *Theoria*, volume 63, 1997. Special issue on non-prioritized belief revision.
6. S. O. Hansson and R. Wassermann. Local change. In *Fourth Symposium on Logical Formalizations of Common Sense Reasoning*, London, UK, 1998.
7. G. Harman. *Change in view – principles of reasoning*. MIT Press, 1986.
8. R. P. Loui. Process and policy: resource-bounded non-demonstrative reasoning. *Computational Intelligence*, 14(1):1–38, 1998.
9. H. Prakken. From logic to dialectics in legal argument. In *Proceedings of the Fifth International Conference on Artificial Intelligence and Law*, pages 165–174, Washington DC, USA, 1995. ACM Press.
10. H. Prakken. Relating protocols for dynamic dispute with logics for defeasible argumentation. *Synthese*, forthcoming, 2000. Special issue on New Perspectives in Dialogical Logic, S. Rahman and H. Rückert, editors.
11. H. Prakken and G. Vreeswijk. Logics for defeasible argumentation. In D. Gabbay, editor, *Handbook of Philosophical Logic (to appear)*. Kluwer Academic Publishers, second edition, 1999.
12. R. Wassermann. Full acceptance through argumentation – a preliminary report. In *Workshop on Practical Reasoning and Rationality at the 16th International Joint Conference on Artificial Intelligence (IJCAI99)*, 1999.
13. R. Wassermann. Resource-bounded belief revision. *Erkenntnis*, 50(2–3), 1999.
14. R. Wassermann. On structured belief bases. In H. Rott and M.-A. Williams, editors, *Frontiers in Belief Revision (to appear)*. Kluwer, 2000.