

Argumentative Credibility-based Revision in Multi-Agent Systems

Luciano H. Tamargo *, Alejandro J. García *,
Matthias Thimm **, and Patrick Krümpelmann ***

* Universidad Nacional del Sur, CONICET, Argentina

** Universität Koblenz-Landau, Germany

*** Technische Universität Dortmund, Germany

Abstract. We consider the problem of belief revision in a multi-agent system with information stemming from different agents with different degrees of credibility. In this context an agent has to carefully choose which information is to be accepted for revision in order to avoid believing in faulty and untrustworthy information. We propose a revision process combining selective revision, deductive argumentation, and credibility information for the adequate handling of information in this complex scenario. New information is evaluated based on the credibility of the source in combination with all arguments favoring and opposing the new information. The evaluation process determines which part of the new information is to be accepted for revision and thereupon incorporated into the belief base by an appropriate revision operator. We demonstrate the benefits of our approach, investigate formal properties, and show that it outperforms the baseline approach without argumentation.

Keywords: Argumentation, Belief Revision, Multi-agent System

1 Introduction

Revising an agent's beliefs is a crucial operation when the agent is situated in a changing environment and only incomplete information is at hand. The area of belief revision [1] is concerned with revising the beliefs of a single agent in the light of new information. The focus of this area is on *prioritized revision*, i. e., new information takes precedence over current beliefs and new information is always accepted. In a dynamic environment with multiple agents this approach is, in general, not apt as information stemming from some agent might be wrong due to unawareness, lack of competence, or even by intention. In order to cope with this situation one has also to take the *credibility* (or *trust*) of the sources into account when accepting new information [2, 3].

In this paper we consider the problem of revising an agent's beliefs with information coming from different agents in a multi-agent system using credibility. We build on previous work on multi-agent revision [2] and employ argumentation strategies for the actual revision process [4]. Computational argumentation [5] is a default reasoning technique that uses arguments and counterarguments to infer those pieces of information from a set of beliefs that are *warranted*, i. e., pieces of information with reasonable grounds to believe. For the purpose of deciding whether new information can be accepted for revision we employ *deductive*

argumentation [6, 5]. The framework of deductive argumentation is a specific approach to structured argumentation that—opposed to abstract approaches such as [7]—uses propositional logic for knowledge representation. Then arguments for formulas are simply proofs and inference is realized by comparing arguments with contradicting counterarguments in *argument trees*.

Here, we extend the framework of deductive argumentation to consider the credibility of the sources of information within the argumentation process. Further, we embed this argumentation framework into a *selective revision* process [8] and obtain a procedure that uses argumentation for deciding whether and which new pieces of information should be accepted for revision. In particular, we consider multiple belief base revision in which the information to be revised by consists of a set of formulas instead of a single formula. Consequently each formula can be evaluated and accepted separately. However, the set of formulas can jointly form an argument for some formula or can support each other. The procedure we propose considers all constructible arguments, evaluates their respective credibility and determines the acceptance status of each new formula. The actual multiple base revision is then performed by the set of accepted formulas of the new information. We investigate the benefits of using argumentation for belief revision in multi-agent systems and show that our approach outperforms the baseline approach without argumentation.

The rest of this paper is organized as follows. In the next section we introduce necessary preliminaries for our investigation. This comprises of logical background, a brief overview on belief revision, and an introduction to deductive argumentation. We continue with presenting an epistemic model based on credibility for an agent situated in a multi-agent environment. Afterwards we present our approach to argumentative credibility-based revision of epistemic models and go on with a throughout analysis of this approach. Finally, we review related work and conclude with a conclusion.

2 Preliminaries

The beliefs of an agent are given in the form of propositional formulas. Let \mathcal{L} be a propositional language generated by some set of atoms and the connectives \wedge , \vee , \Rightarrow , and \neg . As a notational convenience we assume some arbitrary total order \ll on the elements of \mathcal{L} which is used to enumerate elements of each finite $\Phi \subseteq \mathcal{L}$ in a unique way, cf. [6]. For a finite subset $\Phi \subseteq \mathcal{L}$ the *canonical enumeration* of Φ is the vector $\langle \phi_1, \dots, \phi_n \rangle$ such that $\{\phi_1, \dots, \phi_n\} = \Phi$ and $\phi_i \ll \phi_j$ for every $i < j$ with $i, j = 1, \dots, n$. As \ll is total the canonical enumeration of every finite subset $\Phi \subseteq \mathcal{L}$ is uniquely defined.

We use the operator \vdash to denote classical entailment, i.e., for $\Phi_1, \Phi_2 \subseteq \mathcal{L}$ we write $\Phi_1 \vdash \Phi_2$ if and only if Φ_2 is classically entailed by Φ_1 . For $\phi, \phi' \in \mathcal{L}$ we write $\phi \vdash \phi'$ instead of $\{\phi\} \vdash \{\phi'\}$. The deductive closure $\text{Cn}(\Phi) \subseteq \mathcal{L}$ of Φ is defined as $\text{Cn}(\Phi) = \{\phi \in \mathcal{L} \mid \Phi \vdash \phi\}$. Two sets of formulas $\Phi, \Phi' \subseteq \mathcal{L}$ are *equivalent*, denoted by $\Phi \equiv^p \Phi'$, if and only if $\Phi \vdash \Phi'$ and $\Phi' \vdash \Phi$. We also use the equivalence relation \cong^p which is defined as $\Phi \cong^p \Phi'$ if and only if there is a bijection $\sigma : \Phi \rightarrow \Phi'$ such that for every $\phi \in \Phi$ it holds that $\phi \equiv^p \sigma(\phi)$. This means that $\Phi \cong^p \Phi'$ if Φ and Φ' are *element-wise* equivalent. Note that $\Phi \cong^p \Phi'$ implies $\Phi \equiv^p \Phi'$ but not vice

versa. In particular, it holds that e. g. $\{a \wedge b\} \equiv^p \{a, b\}$ but $\{a \wedge b\} \not\equiv^p \{a, b\}$. For $\phi, \phi' \in \mathcal{L}$ we write $\phi \equiv \phi'$ instead of $\{\phi\} \equiv \{\phi'\}$ if $\equiv \in \{\equiv^p, \cong^p\}$. If $\Phi \vdash \perp$ we say that Φ is *inconsistent*. For a set S let $\mathfrak{P}(S)$ denote the power set of S , i.e. the set of all subsets of S . For a set S let $\mathfrak{P}\mathfrak{P}(S)$ denote the set of multi-sets of S , i.e. the set of all subsets of S where an element may occur more than once. To distinguish sets from multi-sets we use brackets “ $\langle \rangle$ ” and “ $\{ \}$ ” for the latter.

Belief Revision. The field of belief revision is concerned with the change of beliefs when more recent or more reliable information is at hand [9, 1]. In this paper, we consider the problem of *multiple belief base revision* [4], cf. the notions of *multiple revision* [1] and *parallel belief revision* [10]. That is, given a finite set of formulas $\mathcal{K} \subseteq \mathcal{L}$ (the *belief base*) and another finite set of formulas $\Phi \subseteq \mathcal{L}$ (the new information) we are interested in revision operations of the form $\mathcal{K} * \Phi$. We distinguish between *prioritized revision*—which requires $\mathcal{K} * \Phi \vdash \Phi$ to hold—and *non-prioritized revision*—which does not necessarily require $\mathcal{K} * \Phi \vdash \Phi$ to hold. Let $\mathcal{K} + \Phi$ be the standard expansion defined via $\mathcal{K} \cup \Phi$. Some rationality postulates apt for multiple base revision can be phrased as follows [4].

Success. $\mathcal{K} * \Phi \vdash \Phi$.

Inclusion. $\mathcal{K} * \Phi \subseteq \mathcal{K} + \Phi$.

Vacuity. If $\mathcal{K} \cup \Phi \not\vdash \perp$ then $\mathcal{K} + \Phi \subseteq \mathcal{K} * \Phi$.

Consistency. If Φ is consistent then $\mathcal{K} * \Phi$ is consistent.

Relevance. If $\alpha \in (\mathcal{K} \cup \Phi) \setminus (\mathcal{K} * \Phi)$ then there is a set H such that $\mathcal{K} * \Phi \subseteq H \subseteq \mathcal{K} \cup \Phi$ and H is consistent but $H \cup \{\alpha\}$ is inconsistent.

Weak Extensionality. If $\Phi \cong^p \Phi'$ then $\mathcal{K} * \Phi \equiv^p \mathcal{K} * \Phi'$.

Weak Success. If $\mathcal{K} \cup \Phi \not\vdash \perp$ then $\mathcal{K} * \Phi \vdash \Phi$.

Consistent Expansion. If $\mathcal{K} \not\subseteq \mathcal{K} * \Phi$ then $\mathcal{K} \cup (\mathcal{K} * \Phi) \vdash \perp$.

We say that $*$ is a *prioritized revision operator* if $*$ satisfies *success*, *inclusion*, *vacuity*, *consistency*, *relevance*, and *weak extensionality*. We say that $*$ is a *non-prioritized revision operator* if $*$ satisfies *inclusion*, *consistency*, *weak extensionality*, *weak success*, and *consistent expansion*.

A specific approach to non-prioritized belief revision is *selective revision* [8]. In the spirit of [4] we apply selective revision to the problem of multiple base revision as follows. For finite $\mathcal{K} \subseteq \mathcal{L}$ a *transformation function* $f_{\mathcal{K}}$ is a function $f_{\mathcal{K}} : \mathfrak{P}(\mathcal{L}) \rightarrow \mathfrak{P}(\mathcal{L})$. Consider the following properties [4].

Inclusion. $f_{\mathcal{K}}(\Phi) \subseteq \Phi$

Weak Extensionality. If $\Phi \cong^p \Phi'$ then $f_{\mathcal{K}}(\Phi) \cong^p f_{\mathcal{K}}(\Phi')$

Consistency Preservation. If Φ is consistent then $f_{\mathcal{K}}(\Phi)$ is consistent

Weak Maximality. If $\mathcal{K} \cup \Phi$ is consistent then $f_{\mathcal{K}}(\Phi) = \Phi$

Using transformation functions we can establish a relationship between prioritized and non-prioritized revision operators as follows [4].

Proposition 1. *Let $*$ be a prioritized revision operator and let $f_{\mathcal{K}}$ satisfy inclusion, weak extensionality, consistency preservation, and weak maximality. Then \circ defined via $\mathcal{K} \circ \Phi = \mathcal{K} * f_{\mathcal{K}}(\Phi)$ for finite $\mathcal{K}, \Phi \subseteq \mathcal{L}$ is a non-prioritized multiple base revision operator.*

Deductive Argumentation. Argumentation frameworks [11] allow for reasoning with inconsistent information based on the notions of arguments, counterarguments and their relationships. In this paper we use the framework of *deductive argumentation* as proposed by Besnard and Hunter [6].

Definition 1. An argument \mathcal{A} for $\phi \in \mathcal{L}$ in $\Phi \subseteq \mathcal{L}$ is a tuple $\mathcal{A} = \langle \Psi, \phi \rangle$ with $\Psi \subseteq \Phi$ such that (1) $\Psi \not\vdash \perp$, (2) $\Psi \vdash \phi$, and (3) there is no $\Psi' \subsetneq \Psi$ with $\Psi' \vdash \phi$.

Hence, an argument $\mathcal{A} = \langle \Psi, \phi \rangle$ for ϕ is a minimal proof for entailing ϕ . For $\mathcal{A} = \langle \Psi, \phi \rangle$ we define $\text{clai}(\mathcal{A}) = \phi$ and $\text{support}(\mathcal{A}) = \Psi$. An argument $\mathcal{A} = \langle \Psi, \phi \rangle$ is more *conservative* than an argument $\mathcal{B} = \langle \Phi, \phi' \rangle$ if and only if $\Psi \subseteq \Phi$ and $\phi' \vdash \phi$. An argument \mathcal{A} is *strictly more conservative* than an argument \mathcal{B} if and only if \mathcal{A} is more conservative than \mathcal{B} but \mathcal{B} is not more conservative than \mathcal{A} .

Definition 2. An argument $\mathcal{A} = \langle \Psi, \phi \rangle$ is an undercut for an argument $\mathcal{B} = \langle \Phi, \phi' \rangle$ if and only if $\phi = \neg(\phi_1 \wedge \dots \wedge \phi_n)$ for some $\{\phi_1, \dots, \phi_n\} \subseteq \Phi$.

If \mathcal{A} is an undercut for \mathcal{B} then we also say that \mathcal{A} *attacks* \mathcal{B} . An argument $\mathcal{A} = \langle \Psi, \phi \rangle$ is a *maximally conservative undercut* for an argument $\mathcal{B} = \langle \Phi, \phi' \rangle$ if and only if \mathcal{A} is an undercut of \mathcal{B} and there is no undercut \mathcal{A}' for \mathcal{B} that is strictly more conservative than \mathcal{A} . An argument $\mathcal{A} = \langle \Psi, \neg(\phi_1 \wedge \dots \wedge \phi_n) \rangle$ is a *canonical undercut* for an argument $\mathcal{B} = \langle \Phi, \phi \rangle$ if and only if \mathcal{A} is a maximally conservative undercut for \mathcal{B} and $\langle \phi_1, \dots, \phi_n \rangle$ is the canonical enumeration of Φ .

It can be shown that it suffices to consider only the canonical undercuts for an argument in order to come up with a reasonable argumentative evaluation of some claim ϕ [6].

Definition 3. Let $\phi \in \mathcal{L}$ be some sentence and let $\Phi \subseteq \mathcal{L}$. An argument tree $\tau_\Phi(\phi)$ for ϕ in Φ is a tree where the nodes are arguments and that satisfies (1) the root is an argument for ϕ in Φ , (2) for every path $[\langle \Phi_1, \phi_1 \rangle, \dots, \langle \Phi_n, \phi_n \rangle]$ in $\tau_\Phi(\phi)$ it holds that $\Phi_n \not\subseteq \Phi_1 \cup \dots \cup \Phi_{n-1}$, and (3) the children $\mathcal{B}_1, \dots, \mathcal{B}_m$ of a node \mathcal{A} consist of all canonical undercuts for \mathcal{A} that obey 2.). Let $\mathcal{T}(\mathcal{L})$ be the set of all argument trees.

An argument tree is a concise representation of the relationships between different arguments that favor or reject some argument \mathcal{A} . In order to evaluate whether a claim ϕ can be justified we have to consider all argument trees for ϕ and all argument trees for $\neg\phi$. For an argument tree τ let $\text{root}(\tau)$ denote the root node of τ . Furthermore, for a node $\mathcal{A} \in \tau$ let $\text{ch}_\tau(\mathcal{A})$ denote the children of \mathcal{A} in τ and let $\text{ch}_\tau^{\mathcal{T}}(\mathcal{A})$ denote the set of sub-trees rooted at a child of \mathcal{A} .

Definition 4. Let $\phi \in \mathcal{L}$ and let $\Phi \subseteq \mathcal{L}$. The argument structure $\Gamma_\Phi(\phi)$ for ϕ wrt. Φ is the tuple $\Gamma_\Phi(\phi) = (\mathcal{P}, \mathcal{C})$ such that \mathcal{P} is the set of argument trees for ϕ in Φ and \mathcal{C} is the set of arguments trees for $\neg\phi$ in Φ .

The argument structure $\Gamma_\Phi(\phi)$ of $\phi \in \mathcal{L}$ gives a complete picture of the reasons for and against ϕ . We use argument structures later to implement a procedure that decides whether some given piece of information should be accepted or rejected for revision.

3 Credibility-based Epistemic Models

We continue with developing an epistemic model for an agent in a multi-agent environment that takes credibilities of other agents into account. Our formalization is based on [2]. Let $\mathbb{A} = \{A_1, \dots, A_n\}$ be a finite set of agents.

Definition 5. If $\phi \in \mathcal{L}$ and $A \in \mathbb{A}$ then $A:\phi$ is called an information object. Let $\mathfrak{I}(\mathcal{L}, \mathbb{A})$ denote the set of all information objects wrt. \mathcal{L} and \mathbb{A} .

An information object $A:\phi$ states that ϕ has been uttered by A . For $\mathcal{I} \subseteq \mathfrak{I}(\mathcal{L}, \mathbb{A})$ we abbreviate $\text{Form}(\mathcal{I}) = \{\phi \mid A:\phi \in \mathcal{I}\}$. We extend the operator $\text{Cn}(\cdot)$ to $\mathfrak{I}(\mathcal{L}, \mathbb{A})$ by defining $\text{Cn}(\mathcal{I}) = \text{Cn}(\text{Form}(\mathcal{I}))$. Note that we do not consider nested information objects such as “ A said that A' said that ϕ ” to keep things simple. We leave this issue for future work.

Remark 1. Although the framework of deductive argumentation from the previous section has been phrased for the language \mathcal{L} we adopt the notions in the same manner for $\mathfrak{I}(\mathcal{L}, \mathbb{A})$ by ignoring the annotated sources. For example, if $\mathcal{I} \subseteq \mathfrak{I}(\mathcal{L}, \mathbb{A})$ and $A:\phi \in \mathfrak{I}(\mathcal{L}, \mathbb{A})$ then we say that $\langle \mathcal{I}, \phi \rangle$ is an argument whenever $\langle \text{Form}(\mathcal{I}), \phi \rangle$ is an argument.

For $\mathcal{I} \subseteq \mathfrak{I}(\mathcal{L}, \mathbb{A})$ with $\text{Form}(\mathcal{I}) \not\perp$ and a total preorder \leq on \mathbb{A} (called *credibility order*) the tuple (\mathcal{I}, \leq) is called a *belief base*. If $\mathcal{K}_A = (\mathcal{I}_A, \leq_A)$ is the belief base of an agent A then $A' \leq_A A''$ means that A believes that A'' is at least as credible as A' . The strict relation $<_A$ and the equivalence relation \equiv_A are defined as usual.

Example 1. Let $\mathbb{A} = \{A_1, A_2, A_3\}$ be a set of agents and consider the belief base $\mathcal{K}_{A_1} = (\mathcal{I}_{A_1}, \leq_{A_1})$ of agent A_1 given via $\mathcal{I}_{A_1} = \{A_1:\neg b, A_2:a, A_3:a \Rightarrow \neg b, A_3:c\}$ and $\leq_{A_1} = A_1 \leq_{A_1} A_2 \leq_{A_1} A_3$. Observe that according to \mathcal{K}_{A_1} , A_1 believes that c has been uttered by A_3 . Furthermore, A_1 believes that A_2 is less credible than A_3 and that himself is less credible than A_2 .

Let $A \in \mathbb{A}$ be an agent and let $\mathcal{K}_A = (\mathcal{I}_A, \leq_A)$ be its belief base. The credibility order \leq_A can be used to specify a preference relation among arguments. Let $\langle \mathcal{I}_1, \phi_1 \rangle, \langle \mathcal{I}_2, \phi_2 \rangle$ be two arguments with $\mathcal{I}_1, \mathcal{I}_2 \subseteq \mathfrak{I}(\mathcal{L}, \mathbb{A})$. Then $\langle \mathcal{I}_1, \phi_1 \rangle$ is *as least as preferred* as $\langle \mathcal{I}_2, \phi_2 \rangle$ by A , denoted by $\langle \mathcal{I}_2, \phi_2 \rangle \preceq_A \langle \mathcal{I}_1, \phi_1 \rangle$ if and only if for all $B:\phi \in \mathcal{I}_1$ there is a $B':\phi' \in \mathcal{I}_2$ such that $B' \leq_A B$. In other words, it holds $\langle \mathcal{I}_2, \phi_2 \rangle \preceq_A \langle \mathcal{I}_1, \phi_1 \rangle$ if and only if the least credible source in \mathcal{I}_1 is at least as credible as the least credible source of \mathcal{I}_2 .

Example 2. Consider \mathcal{K}_{A_1} of Example 1. Let $\langle \mathcal{I}_1, \neg b \rangle$ and $\langle \mathcal{I}_2, c \rangle$ be two arguments with $\mathcal{I}_1 = \{A_2:a, A_3:a \Rightarrow \neg b\}$ and $\mathcal{I}_2 = \{A_3:c\}$. According to \leq_{A_1} , A_2 is less credible than A_3 ($A_2 \leq_{A_1} A_3$) hence $\langle \mathcal{I}_1, \neg b \rangle \preceq_{A_1} \langle \mathcal{I}_2, c \rangle$.

4 Argumentative Credibility-based Revision

Consider a multi-agent system with agents $\mathbb{A} = \{A_1, \dots, A_n\}$ where each agent A_i ($i = 1, \dots, n$) maintains its own belief base $\mathcal{K}_{A_i} = (\mathcal{I}_{A_i}, \leq_{A_i})$. That is, each agent has some subjective beliefs consisting of individual pieces of information


annotated with the source of this information (possibly the agent itself) and some subjective ordering on the credibility of the agents in the system (including itself). When an agent A_j sends some pieces of information $\mathcal{I} \subseteq \mathcal{I}_{A_j}$ to some agent A_i the agent A_i has to deliberate on how to react to receiving \mathcal{I} . Clearly, A_i should not blindly—i.e. in a prioritized fashion—revise \mathcal{I}_{A_i} by \mathcal{I} but take into account the credibility of A_j wrt. \leq_{A_i} . Furthermore, as \mathcal{I} may contain an information object $A_k:\phi$ with $A_k \neq A_j$, i.e. agent A_j forwards some information from A_k to A_i , agent A_i should also consider the credibility of A_k .

Our approach follows the ideas of *selective revision by deductive argumentation* [4] but also incorporates the role of credibilities. On receiving some pieces of information $\mathcal{I} \subseteq \mathcal{I}_{A_j}$ from some agent A_j agent A_i evaluates each $A:\phi \in \mathcal{I}$ by an argumentation procedure that results in either accepting or rejecting $A:\phi$ for revision. This argumentation procedure is regulated by agent A_i 's assessment of the credibilities of the sources of information. In particular, information that comes from a more credible source is preferred to information that comes from a less credible source. A central tool for evaluation in deductive argumentation is a *categorizer*, cf. [6]. A categorizer is meant to assign a value to an argument tree τ depending on how strongly this argument tree favors the root argument. In particular, the larger the value of $\gamma(\tau)$ the better the justification in believing in the claim of the root argument. Here we implement a categorizer as follows.

Definition 6. Let $\mathcal{K}_A = (\mathcal{I}_A, \leq_A)$ be the belief base of an agent A , let $\mathcal{I} \subseteq \mathfrak{J}(\mathcal{L}, \mathbb{A})$, and let τ be some argument tree for $A':\phi$ in \mathcal{I} . Then define the credibility categorizer γ_A^c for A through $\gamma_A^c(\tau) = 1$ if $\text{ch}_\tau(\text{root}(\tau)) = \emptyset$ and through $\gamma_A^c(\tau) = 1 - \max\{\gamma_A^c(\tau') \mid \tau' \in \text{ch}_\tau^T(\text{root}(\tau)) \text{ and } \text{root}(\tau) \preceq_A \text{root}(\tau')\}$ otherwise.

Note that the credibility categorizer implements a similar behavior like grounding semantics for abstract argumentation, cf. [7]. In particular, an argument tree consisting of a single node is categorized with value 1. An argument tree with multiple nodes is categorized with 0 if there is at least one sub-tree under the root that is categorized with 1; and it is categorized with value 1 if all sub-trees under the root are categorized 0. The credibility categorizer takes the subjective credibility order of agent A into account by only considering those sub-trees of a node where the root argument is at least as preferred as the node itself.

Example 3. Let $\mathbb{A} = \{A_1, A_2, A_3\}$ be a set of agents and consider the belief base $\mathcal{K}_{A_1} = (\mathcal{I}_{A_1}, \leq_{A_1})$ of agent A_1 where $\mathcal{I}_{A_1} = \{A_2:b, A_3:c\}$ and $\leq_{A_1} = A_1 \leq_{A_1} A_2 \leq_{A_1} A_3$. Let $\mathcal{I} = \{A_3:a \Rightarrow \neg b, A_2:a\}$. Note that there is exactly one argument tree τ_1 for $a \Rightarrow \neg b$ and one argument tree τ_2 for $a \wedge b$ in $\mathcal{I}_{A_1} \cup \mathcal{I}$. In τ_1 the root is the argument $\mathcal{A} = \langle \{A_3:a \Rightarrow \neg b\}, a \Rightarrow \neg b \rangle$ which has the single canonical undercut $\mathcal{B} = \langle \{A_2:a, A_2:b\}, a \wedge b \rangle$. In τ_2 the situation is reversed and the root of τ_2 is the argument \mathcal{B} which has the single canonical undercut \mathcal{A} . Therefore, the argument structure for $a \Rightarrow \neg b$ is given via $\Gamma_{\mathcal{I}_{A_1} \cup \mathcal{I}}(a \Rightarrow \neg b) = (\{\tau_1\}, \{\tau_2\})$. We can see these argument trees in Figure 1. In τ_1 one can see that the only child of \mathcal{A} is not considered when evaluating with $\gamma_{A_1}^c$ because A_2 is less credible than A_3 according to A_1 . For this reason $\gamma_{A_1}^c(\tau_1) = 1$. However, in τ_2 the situation is reversed and \mathcal{B} is considered by $\gamma_{A_1}^c$. For this reason $\gamma_{A_1}^c(\tau_2) = 0$.

**Fig. 1.** Argument trees in Example 3

We use the credibility categorizer to evaluate new information $\mathcal{I} \subseteq \mathcal{I}_{A_j}$ received by an agent A_i from agent A_j on an argumentative basis and by taking credibilities into account. We say that an agent A_i with belief base $\mathcal{K}_{A_i} = (\mathcal{I}_{A_i}, \leq_{A_i})$ *credulously accepts* an information object $A:\phi \in \mathcal{I}$ wrt. \mathcal{I} if and only if

$$\kappa^c(\mathcal{P}, \mathcal{C}) = \sum_{\tau \in \mathcal{P}} \gamma_{A_i}^c(\tau) - \sum_{\tau \in \mathcal{C}} \gamma_{A_i}^c(\tau) \geq 0 \quad (1)$$

where $\Gamma_{\mathcal{I}_{A_i} \cup \mathcal{I}}(A:\phi) = (\mathcal{P}, \mathcal{C})$ is the argument structure for $A:\phi$ wrt. $\mathcal{I}_{A_i} \cup \mathcal{I}$. The function κ^c is called the *simple accumulator*, cf. [4]. Equation (1) means that A_i accepts $A:\phi$ if there are at least as many reasons to believe in ϕ as there are to believe in $\neg\phi$. The agent A_i *skeptically accepts* $A:\phi \in \mathcal{I}$ wrt. \mathcal{I} if and only if $\kappa^c(\mathcal{P}, \mathcal{C}) > 0$. Using the notion of acceptance we define transformation functions C_{A_i} and S_{A_i} for agent A_i via $C_{A_i}(\mathcal{I}) = \{A:\phi \in \mathcal{I} \mid A_i \text{ cred. accepts } A:\phi \text{ wrt. } \mathcal{I}\}$ and $S_{A_i}(\mathcal{I}) = \{A:\phi \in \mathcal{I} \mid A_i \text{ skept. accepts } A:\phi \text{ wrt. } \mathcal{I}\}$.

Note that—in contrast to the transformation functions discussed before—the codomains of C_{A_i} and S_{A_i} are subsets of $\mathfrak{J}(\mathcal{L}, \mathbb{A})$ instead of \mathcal{L} . We now turn to the issue of revising \mathcal{I}_{A_i} in a prioritized fashion by $C_{A_i}(\mathcal{I})$ and $S_{A_i}(\mathcal{I})$, respectively. We do this by exploiting the Levi-identity for belief revision [9], i.e. by first contracting \mathcal{I}_{A_i} by the complement of $C_{A_i}(\mathcal{I})$ ($S_{A_i}(\mathcal{I})$) and then expanding by $C_{A_i}(\mathcal{I})$ ($S_{A_i}(\mathcal{I})$). Let $-$ be some belief base contraction—for instance a kernel contraction [1]—and define a contraction $-_b$ on $\mathfrak{J}(\mathcal{L}, \mathbb{A})$ for finite $\mathcal{I} \in \mathfrak{J}(\mathcal{L}, \mathbb{A})$ and $\phi \in \mathcal{L}$ through $\mathcal{I} -_b \phi = \{A:\phi' \in \mathcal{I} \mid \phi' \in \text{Form}(\mathcal{I}) - \phi\}$.

Then, for finite $\mathcal{I}, \mathcal{I}' \in \mathfrak{J}(\mathcal{L}, \mathbb{A})$ with $\text{Form}(\mathcal{I}) \not\perp$ define a (prioritized) revision $*$ through

$$\mathcal{I} * \mathcal{I}' = (\mathcal{I} -_b \bigvee_{\phi \in \text{Form}(\mathcal{I}')} \neg\phi) \cup \mathcal{I}' \quad (2)$$

and (non-prioritized) revisions \circ_A^C and \circ_A^S wrt. an agent A through $\mathcal{I} \circ_A^C \mathcal{I}' = \mathcal{I} * C_A(\mathcal{I}')$ and $\mathcal{I} \circ_A^S \mathcal{I}' = \mathcal{I} * S_A(\mathcal{I}')$.

As we stated above, to revise \mathcal{I}_{A_i} for a set of information objects \mathcal{I} , we should contract \mathcal{I}_{A_i} by the complement of \mathcal{I} . For a given set of information objects \mathcal{I} where $\text{Form}(\mathcal{I}) = \{\phi_1, \dots, \phi_n\}$, $\bigvee_{\phi \in \text{Form}(\mathcal{I})} \neg\phi \vdash \{\neg\phi_1, \dots, \neg\phi_n\}$. However, defining the complement of \mathcal{I} as $\{\neg\phi_1, \dots, \neg\phi_n\}$ and using a multiple contraction operator as in [12] would not be sufficient as the following example illustrates.

Example 4. Assume $\mathcal{I}_{A_i} = \{\neg a \vee \neg b\}$ and $\mathcal{I} = \{a, b\}$. Any reasonable contraction operator, cf. [12], would change \mathcal{I}_{A_i} in a minimal way such that $\mathcal{I}_{A_i} - \{\neg a, \neg b\} \not\vdash \neg a$ and $\mathcal{I}_{A_i} - \{\neg a, \neg b\} \not\vdash \neg b$. In this case we get $\mathcal{I}_{A_i} - \{\neg a, \neg b\} = \mathcal{I}_{A_i}$, but obviously $\mathcal{I}_{A_i} \cup \mathcal{I} \vdash \perp$.

5 Analysis

We first illustrate our approach with an example.

Example 5. Imagine the agent *Anna* wants to spend her holidays on Hawaii. Anna's boss *Bob* doesn't want Anna to go on vacation at this time of the year and tells her that she has to do some work. However, Anna is aware of the fact that *Paul*, a good colleague, can do her work. Now Paul gets ill—and therefore cannot take Anna's duties—so Anna has to revise her beliefs accordingly.

In this scenario let $\mathbb{A} = \{A_a, A_p, A_b, A_c\}$ where A_a is Anna, A_p is an Anna's colleague Paul, A_b is an Anna's boss, and A_c is the only client of the company. Consider the sentences: *Anna travels to Hawaii* (abbreviated h), *there is work to do* (w), *Paul can do Anna's work* (r), and *Paul is ill* (i).

Now consider Anna's belief base \mathcal{K}_{A_a} given via $\mathcal{I}_{A_a} = \{A_a : h, A_c : \neg w, A_p : r, A_p : r \Rightarrow h, A_a : \neg w \Rightarrow h\}$. This means that Anna believes that she should travel to Hawaii ($A_a : h$), that there is no work to do because the only client (A_c) of the company said so ($A_c : \neg w$), that Paul can fill in for her ($A_p : r$), that if she has a replacement then she can go to Hawaii $A_p : r \Rightarrow h$, and that if there is no work to do then she can go to Hawaii $A_a : \neg w \Rightarrow h$. Furthermore, let the credibility order among agents according to Anna \leq_{A_a} be defined via $A_a \leq_{A_a} A_p \leq_{A_a} A_b \leq_{A_a} A_c$.

Now consider the new information $\Phi = \{A_b : w, A_p : i, A_b : i \Rightarrow \neg r\}$ stemming from communication with Anna's boss which states that there is work to do ($A_b : w$), that Paul is ill ($A_p : i$), and that if Paul is ill then Anna has no replacement ($A_b : i \Rightarrow \neg r$). As one can see there are some arguments for and against w , i and r in $\mathcal{I}_{A_a} \cup \Phi$, e. g., arguments for and against w are $\langle \{A_b : w\}, w \rangle$, $\langle \{A_c : \neg w\}, \neg w \rangle$.

We compute the argument structures $\Gamma_{\mathcal{I}_{A_a} \cup \Phi}(\alpha) = (\mathcal{P}, \mathcal{C})$ for each sentence $\alpha \in \text{Form}(\Phi)$ with respect to $\mathcal{I}_{A_a} \cup \Phi$ as follows.

- (w). There is exactly one argument tree τ_1 for w and one argument tree τ_2 for $\neg w$ in $\mathcal{I}_{A_a} \cup \Phi$. In τ_1 the root is the argument $\mathcal{A} = \langle \{A_b : w\}, w \rangle$ which has the single canonical undercut $\mathcal{B} = \langle \{A_c : \neg w\}, \neg w \rangle$. In τ_2 the situation is reversed and the root of τ_2 is the argument \mathcal{B} which has the single canonical undercut \mathcal{A} . Therefore, the argument structure for w is given via $\Gamma_{\mathcal{I}_{A_a} \cup \Phi}(w) = (\{\tau_1\}, \{\tau_2\})$. It follows that $\gamma_{A_a}^c(\tau_1) = 0$, $\gamma_{A_a}^c(\tau_2) = 1$ and $\sum_{\tau \in \mathcal{P}} \gamma_{A_a}^c(\tau_1) - \sum_{\tau \in \mathcal{C}} \gamma_{A_a}^c(\tau_2) = -1$ that $\mathcal{I}_{A_a} \cup \Phi$ which means that w is *rejected*.
- (i). There is exactly one argument tree τ_1 for i and one argument tree τ_2 for $\neg i$ in $\mathcal{I}_{A_a} \cup \Phi$. In τ_1 the root is the argument $\mathcal{A} = \langle \{A_p : i\}, i \rangle$ which has the single canonical undercut $\mathcal{B} = \langle \{A_b : i \Rightarrow \neg r, A_p : r\}, \neg i \rangle$. In τ_2 the situation is reversed and the root of τ_2 is the argument \mathcal{B} which has the single canonical undercut \mathcal{A} . Therefore, the argument structure for i is given via $\Gamma_{\mathcal{I}_{A_a} \cup \Phi}(i) = (\{\tau_1\}, \{\tau_2\})$. It follows that $\gamma_{A_a}^c(\tau_1) = \gamma_{A_a}^c(\tau_2) = 0$ and $\sum_{\tau \in \mathcal{P}} \gamma_{A_a}^c(\tau_1) - \sum_{\tau \in \mathcal{C}} \gamma_{A_a}^c(\tau_2) = 0$ which means that the status of i is *undecided*.
- ($i \Rightarrow \neg r$). There is exactly one argument tree τ_1 for $i \Rightarrow \neg r$ and one argument tree τ_2 for $i \wedge r$ in $\mathcal{I}_{A_a} \cup \Phi$. In τ_1 the root is the argument $\mathcal{A} = \langle \{A_b : i \Rightarrow$

$\neg r\}$, $i \Rightarrow \neg r\}$ which has the single canonical undercut $\mathcal{B} = \langle \{A_p:i, A_p:r\}, i \wedge r \rangle$. In τ_2 the situation is reversed and the root of τ_2 is the argument \mathcal{B} which has the single canonical undercut \mathcal{A} . Therefore, the argument structure for $i \Rightarrow \neg r$ is given via $\Gamma_{\mathcal{I}_{A_a} \cup \Phi}(i \Rightarrow \neg r) = (\{\tau_1\}, \{\tau_2\})$. It follows that $\gamma_{A_a}^c(\tau_1) = 1$, $\gamma_{A_a}^c(\tau_2) = 0$ and $\sum_{\tau \in \mathcal{P}} \gamma_{A_a}^c(\tau_1) - \sum_{\tau \in \mathcal{C}} \gamma_{A_a}^c(\tau_2) = 1$ which means that $i \Rightarrow \neg r$ is *accepted*.

Due to the above evaluation the values of $C_{A_i}(\Phi)$ and $S_{A_i}(\Phi)$ can be determined by $S_{A_a}(\Phi) = \Phi \setminus \{A_b:w, A_p:i\} = \{A_b:i \Rightarrow \neg r\}$ and $C_{A_a}(\Phi) = \Phi \setminus \{A_b:w\} = \{A_p:i, A_b:i \Rightarrow \neg r\}$.

If $*$ is defined via (2) we obtain $\mathcal{I}_{A_a} * S_{A_a}(\Phi) = \{A_a:h, A_c:\neg w, A_p:r, A_p:r \Rightarrow h, A_a:\neg w \Rightarrow h, A_b:i \Rightarrow \neg r\}$ and $\mathcal{I}_{A_a} * C_{A_a}(\Phi) = \{A_a:h, A_c:\neg w, A_p:r \Rightarrow h, A_a:\neg w \Rightarrow h, A_p:i, A_b:i \Rightarrow \neg r\}$.

The above example illustrates that our approach is quite complex and involves a sophisticated deliberation process for deciding how a non-prioritized revision should be performed. One might ask whether the argumentative decision process is necessary and if the same results could be obtained by a simpler approach. For example, consider the following definition of a transformation function.

Definition 7. Let $\mathcal{K}_A = (\mathcal{I}_A, \leq_A)$ and $\mathcal{I} \subseteq \mathfrak{J}(\mathcal{L}, \mathbb{A})$. The function H_A is defined via $H_A(\mathcal{I}) = \{A_i:\phi \in \mathcal{I} \mid \forall \langle \mathcal{I}', \neg\phi \rangle, \mathcal{I}' \subseteq \mathcal{I}_A \cup \mathcal{I}, \langle \mathcal{I}', \neg\phi \rangle \leq_A \langle \{A_i:\phi\}, \phi \rangle\}$.

In other words, the function H_A rejects a $A':\phi \in \mathcal{I}$ if there is a proof for $\neg\phi$ in $\mathcal{I}_A \cup \mathcal{I}$ such that the least credible source of this proof is strictly more credible than A' . Therefore, this definition of a transformation function intuitively implements the idea of how a credibility-based revision should be defined. The question arises whether this definition of a transformation is sufficient for realizing a meaningful revision based on credibilities. In Example 6, we show that this is not the case.

Example 6. Let $\mathbb{A} = \{A_1, A_2, A_3\}$ be a set of agents and consider the belief base $(\mathcal{I}_{A_1}, \leq_{A_1})$ of agent A_1 given via $\mathcal{I}_{A_1} = \{A_3:b, A_3:a \Rightarrow \neg b, A_2:\neg c\}$ and $\leq_{A_1} = A_3 \leq_{A_1} A_2 \leq_{A_1} A_1$. Assume now that A_1 receives the new information \mathcal{I} given via $\mathcal{I} = \{A_3:a \Rightarrow c, A_3:a\}$ and consider the revision of \mathcal{I}_{A_1} by \mathcal{I} . Observe that $C_{A_1}(\mathcal{I}) = \{A_3:a \Rightarrow c\}$ and $H_{A_1}(\mathcal{I}) = \{A_3:a \Rightarrow c, A_3:a\} = \mathcal{I}$. If $*$ is defined via (2) we obtain $\mathcal{I}_{A_1} * C_{A_1}(\mathcal{I}) = \{A_3:b, A_3:a \Rightarrow \neg b, A_2:\neg c, A_3:a \Rightarrow c\}$ and $\mathcal{I}_{A_1} * H_{A_1}(\mathcal{I}) = \{A_3:b, A_3:a \Rightarrow c, A_3:a\}$.

As one can see, the revision based on C_{A_1} differs from the revision based on H_{A_1} which stems from $A_3:a \in H_{A_1}(\mathcal{I})$ and $A_3:a \notin C_{A_1}(\mathcal{I})$. The reason for $A_3:a \in H_{A_1}(\mathcal{I})$ is that there are two proofs for $\neg a$ in $\mathcal{I}_{A_1} \cup \mathcal{I}$ — $\{A_3:b, A_3:a \Rightarrow \neg b\}$ and $\{A_2:\neg c, A_3:a \Rightarrow c\}$ —and the credibility of the least credible agent in both proofs—which is A_3 —is not strictly greater than the credibility of $A_3:a$ —which is A_3 as well. Therefore, H_{A_1} accepts $A_3:a$ for revision. For C_{A_1} the situation is different. As $\langle \{A_3:a\}, a \rangle$ is the only argument for a and there are two arguments— $\langle \{A_3:b, A_3:a \Rightarrow \neg b\}, \neg a \rangle$ and $\langle \{A_2:\neg c, A_3:a \Rightarrow c\}, \neg a \rangle$ —for $\neg a$ the argumentative evaluation of a results in the three argument trees depicted in Figure 2. As all arguments appearing in the argument trees have the same least credible source A_3 no argument is ignored in the evaluation. Therefore

the tree for argument $\langle\{A_3:a\}, a\rangle$ is categorized to 0 and both trees for $\neg a$ are categorized to 1. By (1) it follows that $A_3:a$ is not accepted for revision by C_{A_1} . An implication of this decision is that in $\mathcal{T}_{A_1} * C_{A_1}(\mathcal{T})$ the information $A_2:\neg c$ —which is the single piece of information that comes from more credible information that any other piece of information—is retained.




Fig. 2. Three argument trees of Example 6

As for formal properties for transformation functions and belief revision our approach behaves well. For the following results source annotations of formulas can be neglected.

Proposition 2. *Let A be some agent. The transformation functions S_A and C_A satisfy inclusion, weak inclusion, weak extensionality, consistency preservation, and weak maximality.*

By exploiting Proposition 1, see also [4], we obtain the following result.

Corollary 1. *Let A be some agent. The operators \circ_A^C and \circ_A^S are non-prioritized multiple base revision operators.*

The above corollary shows that argumentative credibility-based revision conforms with expectations to non-prioritized revision.

6 Conclusion and Related Work

To the best of our knowledge there is no other work that uses argumentation for multi-source belief revision with credibilities. However, there are bodies of work on credibility based multi-source belief revision as well as on the use of argumentation in multi-agent systems. Concerning the relation to the former field of research we base our approach on the model for multi-source belief revision of [2] and we use a selective revision operator as presented in [4].

While there has been some work on the revision of argumentation systems, very little work on the application of argumentation techniques for the revision process has been done so far, cf. [13]. In fact, the work most related to the work presented here makes use of negotiation techniques for belief revision [14, 15], without argumentation. In the general setup of [14] a symmetric merging of information from two sources is performed by means of a negotiation procedure that determines which source has to reduce its information in each round. The information to be given up is determined by another function. The negotiation ends when a consistent union of information is reached. While this can be seen as a one step process of merging or consolidation in general, the formalism also

allows to differentiate between the information given up from the first source and the second source. In [14], this setting is then successively biased towards prioritizing the second source which leads to representation theorems for operations equivalent to selective revision satisfying *consistent expansion* and for classic AGM operators. However, the negotiation framework used in [14] is very different from the argumentation formalism used here and also very different from the setup of selective revision. Moreover, the functions for the negotiation and concession are left abstract.

In [15] mutual belief revision is considered where two agents revise their respective belief states by information of the other agent. Both agents agree in a negotiation on the information that is accepted by each agent. The revisions of the agents are split into a selection function and two iterated revision functions which leads to operators satisfying *consistent expansion*. The selection function is then a negotiation function on two sets of beliefs that represent the sets of belief that each agent is willing to accept from the other agent. This setting has a very different focus as ours and also does not specify the selection function.

There is also work on the use of argumentation to reason about trust with [16] being the most recent work in this area. In [16] a meta-argumentation approach is used to argue not only taking the trustworthiness of information sources into account while evaluating the acceptance of arguments, but also to argue about the trustworthiness itself. In these approaches it is determined for a given set of arguments from different sources which ones are accepted in which ones are not. Dynamics of the system in terms of belief revision are not considered. In contrast, here we treat a belief revision problem of non-argumentative belief bases by employing argumentation in the selection process of belief revision.

While the concepts of trust and reputation are complex, in this approach we have taken the position that they can be seen as a kind of credibility value that the agents assign to each other. In contrast to this paper, in [17] a model for reputation is presented that takes into account the social dimension of agents and a hierarchical ontology structure. They show how the model relates to other systems and provide initial experimental results about the benefits of using a social view on the modeling of reputation.

In this paper, we developed a multi-agent revision framework based on using deductive argumentation and credibilities for deciding whether new information should be accepted for revision. We used the very general framework of multiple belief base revision and investigated a scenario where an agent has to revise its belief base of propositional formulas with a set propositional formulas. Formulas are annotated with credibility information of the source of the formula and we developed an argumentation procedure based on credibility that decides which formulas of the set should be accepted for (prioritized) revision. We investigated the properties of our approach and compared it to a simple approach for multi-agent revision and other related work.

Our approach is concerned with revising the actual content of the belief base of an agent given some *static* credibility assessment. That is, the credibilities of the agents in the system are fixed (subjectively) and may not change. How-

ever, this may not be the case in real-world scenarios, see [18] for a discussion. In particular, information received from an agent may change the subjective assessment of its credibility: if an agent often gives good arguments or his information is confirmed by more credible agents then this agent should be assessed more credible as well. The dynamics of credibility assessments can be approached by interpreting credibilities not as annotations but as formulas of the object level themselves and to use traditional revision methods for them as well. Part of future work is on investigating this approach within our framework.

References

1. Hansson, S.O.: A Textbook of Belief Dynamics. Kluwer Academic Publishers, Norwell, MA, USA (2001)
2. Tamargo, L.H., García, A.J., Falappa, M.A., Simari, G.R.: Modeling knowledge dynamics in multi-agent systems based on informants. *The Knowledge Engineering Review (KER)* **27**(1) (2012) 87–114
3. Khosravifar, B., Gomrokchi, M., Bentahar, J., P.Thiran: Maintenance-based trust for open multiagent systems. In: *AAMAS 09, Budapest, Hungary* (2009)
4. Krümpelmann, P., Thimm, M., Falappa, M.A., Garcia, A.J., Kern-Isberner, G., Simari, G.R.: Selective Revision by Deductive Argumentation. In: *1st Int. Workshop on the Theory and Applications of Formal Argumentation (TAFA)*. (2011)
5. Besnard, P., Hunter, A.: *Elements of Argumentation*. The MIT Press (June 2008)
6. Besnard, P., Hunter, A.: A logic-based theory of deductive arguments. *Artificial Intelligence* **128**(1–2) (2001) 203–235
7. Dung, P.M.: On the Acceptability of Arguments and its Fundamental Role in Nonmonotonic Reasoning. *Artificial Intelligence* **77**(2) (1995) 321–358
8. Fermé, E., Hansson, S.O.: Selective revision. *Studia Logica* **63** (1999) 331–342
9. Alchourron, C.E., Gärdenfors, P., Makinson, D.: On the logic of theory change: Partial meet contraction and revision functions. *The Journal of Symbolic Logic* **50**(2) (1985) 510–530
10. Delgrande, J.P., Jin, Y.: Parallel belief revision. In: *Proceedings of the 23rd National Conference on Artificial Intelligence (AAAI-08)*. (2008)
11. Bench-Capon, T.J.M., Dunne, P.E.: Argumentation in artificial intelligence. *Artificial Intelligence* **171** (2007) 619–641
12. Fermé, E., Saez, K., Sanz, P.: Multiple kernel contraction. *Studia Logica: An International Journal for Symbolic Logic* **73**(2) (March 2003) 183–195
13. Falappa, M.A., Kern-Isberner, G., Simari, G.R.: Belief revision and argumentation theory. In: *Argumentation in Artificial Intelligence*. Springer (2009) 341–360
14. Booth, R.: A negotiation-style framework for non-prioritised revision. In: *Proceedings of TARK'01*. (2001) 137–150
15. Zhang, D., Foo, N., Meyer, T., Kwok, R.: Negotiation as mutual belief revision. In: *Proceedings of AAAI'04*. (2004) 317–322
16. Villata, S., Boella, G., Gabbay, D.M., van der Torre, L.: Arguing about the trustworthiness of the information sources. In Liu, W., ed.: *ECSQARU'11*. Volume 6717 of *Lecture Notes in Computer Science*, Springer (2011) 74–85
17. Sabater, J., Sierra, C.: Regret: A reputation model for gregarious societies. *Fourth Workshop on Deception, Fraud and Trust in Agent Societies* (2001) 61–69
18. Tamargo, L.H., Falappa, M.A., García, A.J., Simari, G.R.: A change model for credibility partial order. In: *Proceeding of the 5th International Conference on Scalable Uncertainty Management (SUM)*. (2011) 317–330