



HAL
open science

Graded Reinstatement in Belief Revision

Célia da Costa Pereira, Andrea G. B. Tettamanzi, Serena Villata

► **To cite this version:**

Célia da Costa Pereira, Andrea G. B. Tettamanzi, Serena Villata. Graded Reinstatement in Belief Revision. IEEE/WIC/ACM International Conference on Intelligent Agent Technology, IAT 2011, Aug 2011, Lyon, France. pp.58–61, 10.1109/WI-IAT.2011.185 . hal-01329993

HAL Id: hal-01329993

<https://hal.science/hal-01329993>

Submitted on 9 Jun 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Graded Reinstatement in Belief Revision

Célia da Costa Pereira
Université de Nice Sophia Antipolis
 Nice, France
 celia.pereira@unice.fr

Andrea G. B. Tettamanzi
Università degli Studi di Milano
 Milano, Italy
 andrea.tettamanzi@unimi.it

Serena Villata
INRIA
 Sophia Antipolis, France
 serena.villata@inria.fr

Abstract—An important issue in belief revision is the possible loss of previous information which later might turn to be correct when new information becomes available to the agent. Starting from the fuzzy argument labeling based on trust, we mirror argument reinstatement in belief reinstatement, and we show the conditions under which the belief reinstatement is total, partial, or nonexistent.

Keywords—BDI Agents, Argumentation, Possibility Theory

I. INTRODUCTION

In da Costa Pereira et al. [1], we present a labeling algorithm which assigns to the arguments, proposed by the agents, a fuzzy evaluation. This fuzzy evaluation is originated by the observation that some arguments may come from only partially trusted sources. The arguments then support the agents’ beliefs, represented as the conclusions of structured arguments.

A schematic illustration of the framework proposed in [1] is provided in Figure 1. The framework may be regarded as a belief revision model based on argumentation. An agent interacts with the world by receiving arguments A from one or more *sources*. The agent’s internal mental state is completely described by a fuzzy set of arguments \mathcal{A} , from which the beliefs of the agent may be derived. A *trust* module assigns a trust degree τ to each source. As new arguments A are received, they are added to \mathcal{A} with the same membership degree as the degree τ to which their source is trusted. Fuzzy labeling of \mathcal{A} yields a fuzzy reinstatement labeling α , which may be regarded as a fuzzy set of acceptable arguments, whose consequences induce a possibility distribution π , from which an explicit representation \mathbf{B} of the agent’s beliefs is constructed as the necessity measure N of possibility distribution π .

In this paper, we investigate in details the belief reinstatement proposed in [1]. In particular, we provide and discuss the condition under which a belief reinstatement is total, i.e., the belief previously discarded is reintegrated with a degree $\mathbf{B}(\phi) = 1$, or partial, i.e., the belief is reintegrated with a certain degree $\mathbf{B}(\phi) > 0$, or no reinstatement is allowed. We adopt the representation of uncertain beliefs proposed in [2], [3], and we share the observation of [4], [5] that the “priority to incoming information” presents some drawbacks in multiagent systems where the original belief may be *somehow* reintegrated.

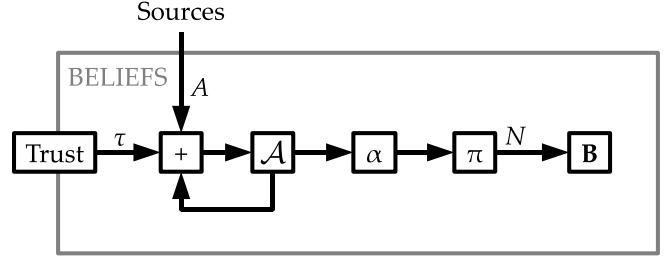


Figure 1. A schematic illustration of the framework [1].

II. BACKGROUND

A Dung’s abstract argumentation [6] is a pair $\langle \mathcal{A}, \rightarrow \rangle$. \mathcal{A} is a set of elements called *arguments* and $\rightarrow \subseteq \mathcal{A} \times \mathcal{A}$ is a binary relation called *attack*. Dung defines a number of acceptability semantics [6].

Dung’s semantics are originally stated in terms of sets of accepted arguments. It is equal to express these concepts using argument *labeling* as proposed by Caminada [7].

Definition 1: (Reinstatement labeling [7]) Let *lab* be an AF-labeling. We say that *lab* is a reinstatement labeling if and only if it satisfies the following:

- $\forall A_i \in \mathcal{A} : (lab(A_i) = out \equiv \exists A_j \in \mathcal{A} : (A_j \rightarrow A_i \wedge lab(A_j) = in))$ and
- $\forall A_i \in \mathcal{A} : (lab(A_i) = in \equiv \forall A_j \in \mathcal{A} : (A_j \rightarrow A_i \supset lab(A_j) = out))$ and
- $\forall A_i \in \mathcal{A} : (lab(A_i) = undec \equiv \exists A_j \in \mathcal{A} : (A_j \rightarrow A_i \wedge \neg(lab(A_j) = out)) \wedge \nexists A_k \in \mathcal{A} : (A_k \rightarrow A_i \wedge lab(A_k) = in))$.

A classical propositional language may be used to represent information for manipulation by a cognitive agent.

Definition 2: (Language) Let Prop be a *finite*. set of atomic propositions and let \mathcal{L} be the propositional language such that $\text{Prop} \cup \{\top, \perp\} \subseteq \mathcal{L}$, and, $\forall \phi, \psi \in \mathcal{L}$, $\neg\phi \in \mathcal{L}$, $\phi \wedge \psi \in \mathcal{L}$, $\phi \vee \psi \in \mathcal{L}$.

As usual, one may define additional logical connectives and consider them as useful shorthands for combinations of connectives of \mathcal{L} , e.g., $\phi \supset \psi \equiv \neg\phi \vee \psi$.

We will denote by $\Omega = \{0, 1\}^{\text{Prop}}$ the set of all interpretations on Prop. An interpretation $\mathcal{I} \in \Omega$ is a function $\mathcal{I} : \text{Prop} \rightarrow \{0, 1\}$ assigning a truth value $p^{\mathcal{I}}$ to every

atomic proposition $p \in \text{Prop}$ and, by extension, a truth value $\phi^{\mathcal{I}}$ to all formulas $\phi \in \mathcal{L}$. We will denote by $[\phi]$ the set of all models of ϕ , $[\phi] = \{\mathcal{I} : \mathcal{I} \models \phi\}$.

We can give arguments a structure, and the attack relation will be defined in terms of such a structure of the arguments.

Definition 3: An argument is a pair $\langle \Phi, \phi \rangle$, with $\phi \in \mathcal{L}$ and $\Phi \subset \mathcal{L}$, such that (i) $\Phi \not\vdash \perp$, (ii) $\Phi \vdash \phi$, (iii) Φ is minimal w.r.t. set inclusion. We call ϕ the conclusion and Φ the support of the argument. An *undercut* for an argument $\langle \Phi, \phi \rangle$ is an argument $\langle \Psi, \psi \rangle$ where $\psi = \neg(\phi_1 \wedge \dots \wedge \phi_n)$ and $\{\phi_1, \dots, \phi_n\} \subseteq \Phi$. A *rebuttal* for an argument $\langle \Phi, \phi \rangle$ is an argument $\langle \Psi, \psi \rangle$ if and only if $\psi \Leftrightarrow \neg\phi$ is a tautology.

III. FUZZY ARGUMENT LABELING

In order to account for the fact that arguments may originate from sources that are trusted only to a certain degree, we [1] extend the (crisp) abstract argumentation structure by allowing gradual membership of arguments in the set of arguments \mathcal{A} . \mathcal{A} is a fuzzy set of arguments, and $\mathcal{A}(A)$, the membership degree of argument A in \mathcal{A} is given by the trust degree of the most reliable (i.e., trusted) source that offers argument A ,

$$\mathcal{A}(A) = \max_{s \in \text{src}(A)} \tau_s, \quad (1)$$

where τ_s is the degree to which source $s \in \text{src}(A)$ is trusted.

This fuzzification of \mathcal{A} provides a natural way of associating strengths to arguments, and suggests rethinking the labeling of an argumentation framework in terms of fuzzy degrees of argument acceptability.

Definition 4: (Fuzzy AF-labeling) Let $\langle \mathcal{A}, \rightarrow \rangle$ be an abstract argumentation framework. A fuzzy AF-labeling is a total function $\alpha : \mathcal{A} \rightarrow [0, 1]$.

Such an α may also be regarded as (the membership function of) the fuzzy set of acceptable arguments: $\alpha(A) = 0$ means the argument is outright unacceptable, $\alpha(A) = 1$ means the argument is fully acceptable, and all cases inbetween are provided for.

Definition 5: (Fuzzy Reinstatement Labeling) Let α be a fuzzy AF-labeling. We say that α is a fuzzy reinstatement labeling if and only if, for all arguments A ,

$$\alpha(A) = \min\{\mathcal{A}(A), 1 - \max_{B: B \rightarrow A} \alpha(B)\}. \quad (2)$$

Such labeling can be computed, by means of an iterative labeling algorithm [1], starting from an initial labeling $\alpha_0 = \mathcal{A}$. Given labeling α_t , the next labeling α_{t+1} is obtained by applying, for all arguments A , the transition step

$$\alpha_{t+1}(A) = \frac{1}{2}\alpha_t(A) + \frac{1}{2}\min\{\mathcal{A}(A), 1 - \max_{B: B \rightarrow A} \alpha_t(B)\}. \quad (3)$$

Equation 3 defines a sequence $\{\alpha_t\}_{t=0,1,\dots}$ of labelings which converges to a fuzzy reinstatement labeling, as it is proved in [1]. Therefore, we define the fuzzy labeling of a fuzzy argumentation framework as the limit of $\{\alpha_t\}_{t=0,1,\dots}$

Definition 6: Let $\langle \mathcal{A}, \rightarrow \rangle$ be a fuzzy argumentation framework. A fuzzy reinstatement labeling for such argumentation framework is, for all argument A ,

$$\alpha(A) = \lim_{t \rightarrow \infty} \alpha_t(A). \quad (4)$$

The convergence speed of the labeling algorithm is linear and a small number of iterations is enough to get so close to the limit that the error is less than the precision with which the membership degrees are represented in the computer.

IV. BELIEFS

In the proposed framework, belief reinstatement is guaranteed thanks to the integration of the argumentation framework with the belief-change phase. More precisely, when a new argument arrives from a partially trusted source, the argumentation framework is updated using the fuzzy labeling algorithm. Therefore, each argument reinstated by the algorithm will induce the reinstatement, to some extent, of the conclusion of the argument in the belief set and of all the formulas that logically follow from the belief set.

As convincingly argued by Dubois and Prade[2], a *belief* should be regarded as a necessity degree induced by a normalized possibility distribution on the set of all interpretations. Starting from such an insight, a fuzzy reinstatement labeling α determines a set of beliefs in a natural way, as we now explain in detail, after reviewing some background on possibility theory.

A. From Arguments to Possibilities

A possibility distribution may be defined as the membership function of a fuzzy set that describes the more or less possible and mutually exclusive values of one (or more) variable(s) [8]. Indeed, if F designates the fuzzy set of possible values of a variable X , $\pi_X = \mu_F$ is called the possibility distribution associated to X . The identity $\mu_F(v) = \pi_X(v)$ means that the membership degree of v to F is equal to the possibility degree of X being equal to v when all we know about X is that its value is in F . A possibility distribution for which there exists a completely possible value ($\exists v_0 : \pi(v_0) = 1$) is said to be *normalized*.

Definition 7: (Possibility and Necessity Measures) A possibility distribution π induces a *possibility measure* and its dual *necessity measure*, denoted by Π and N respectively. Both measures apply to a crisp set A and are defined as follows:

$$\Pi(A) = \sup_{s \in A} \pi(s); \quad (5)$$

$$N(A) = 1 - \Pi(\bar{A}) = \inf_{s \in \bar{A}} \{1 - \pi(s)\}. \quad (6)$$

In words, the possibility measure of set A corresponds to the greatest of the possibilities associated to its elements; conversely, the necessity measure of A is equivalent to the impossibility of its complement \bar{A} .

A few properties of possibility and necessity measures induced by a normalized possibility distribution on a finite universe of discourse Ω are the following. For all subsets $A, B \subseteq \Omega$:

- 1) $\Pi(A \cup B) = \max\{\Pi(A), \Pi(B)\}$;
- 2) $\Pi(\emptyset) = N(\emptyset) = 0$, $\Pi(\Omega) = N(\Omega) = 1$;
- 3) $N(A \cap B) = \min\{N(A), N(B)\}$;
- 4) $\Pi(A) = 1 - N(\bar{A})$ (duality);
- 5) $N(A) \leq \Pi(A)$;
- 6) $N(A) > 0$ implies $\Pi(A) = 1$;
- 7) $\Pi(A) < 1$ implies $N(A) = 0$;

A consequence of these properties is that $\max\{\Pi(A), \Pi(\bar{A})\} = 1$. In case of complete ignorance on A , $\Pi(A) = \Pi(\bar{A}) = 1$.

The beliefs of an agent may thus be completely described by a normalized possibility distribution

$$\pi : \Omega \rightarrow [0, 1], \quad (7)$$

which represents a plausibility order of possible states of affairs: $\pi(\mathcal{I})$ is the possibility degree of interpretation \mathcal{I} .

Given argument $A = \langle \Phi, \phi \rangle$, let $\text{con}(A)$ denote the conclusion of A , i.e., $\text{con}(\langle \Phi, \phi \rangle) = \phi$. The possibility distribution π induced by a fuzzy argumentation framework α may be constructed by letting, for all interpretation \mathcal{I} ,

$$\pi(\mathcal{I}) = \min\{1, 1 + \max_{A:\mathcal{I} \models \text{con}(A)} \alpha(A) - \max_{B:\mathcal{I} \not\models \text{con}(B)} \alpha(B)\}. \quad (8)$$

The first maximum in the above equation accounts for the most convincing argument compatible with world \mathcal{I} , whereas the second maximum accounts for the most convincing argument against world \mathcal{I} . A world will be possible to an extent proportional to the difference between the most convincing argument supporting it and the most convincing argument against it. The world will be considered completely possible if such difference is positive or null, but it will be considered less and less possible (or plausible) as such difference grows more and more negative.

Theorem 1: Any π defined as per Eq. 8 is normalized.

Proof: Either $\pi(\mathcal{I}) = 1$ for all \mathcal{I} , and π is trivially normalized, or there exists an interpretation, say \mathcal{I}_0 , such that $\pi(\mathcal{I}_0) < 1$. By Equation 8, then, it must be

$$\max_{A:\mathcal{I}_0 \models \text{con}(A)} \alpha(A) < \max_{B:\mathcal{I}_0 \not\models \text{con}(B)} \alpha(B).$$

But then, let us consider the complementary interpretation $\bar{\mathcal{I}}_0$, which maps all atoms to a truth value that is the opposite of the truth value they are mapped to by \mathcal{I}_0 . Clearly, all formulas satisfied by \mathcal{I}_0 are not satisfied by $\bar{\mathcal{I}}_0$ and *vice versa*. Therefore,

$$\pi(\bar{\mathcal{I}}_0) = \min\{1, 1 + \max_{B:\bar{\mathcal{I}}_0 \not\models \text{con}(B)} \alpha(B) - \max_{A:\bar{\mathcal{I}}_0 \models \text{con}(A)} \alpha(A)\} = 1.$$

In other words, if a world is not completely plausible, its opposite must be completely plausible, and for this reason π is always normalized. ■

B. Belief Set

The degree to which a given arbitrary formula $\phi \in \mathcal{L}$ is believed can be calculated from the possibility distribution induced by the fuzzy argumentation framework as

$$\mathbf{B}(\phi) = N([\phi]) = 1 - \max_{\mathcal{I} \not\models \phi} \{\pi(\mathcal{I})\}. \quad (9)$$

Such \mathbf{B} may be regarded, at the same time, as a fuzzy modal epistemic operator or as a fuzzy subset of \mathcal{L} .

A powerful feature of such an approach based on a possibility distribution is that $\mathbf{B}(\phi)$ can be computed for any formula ϕ , not just for formulas that are the conclusion of some argument. E.g., if A is an argument whose conclusion is p and B is an argument whose conclusion is $p \supset q$, and $\alpha(A) = \alpha(B) = 1$, then not only $\mathbf{B}(p) = \mathbf{B}(p \supset q) = 1$, but also $\mathbf{B}(q) = 1$, $\mathbf{B}(p \wedge q) = 1$, etc.

Straightforward consequences of the properties of possibility and necessity measures are that $\mathbf{B}(\phi) > 0 \Rightarrow \mathbf{B}(\neg\phi) = 0$, which means that if the agent somehow believes ϕ then it cannot believe $\neg\phi$ at all, and

$$\mathbf{B}(\top) = 1, \quad (10)$$

$$\mathbf{B}(\perp) = 0, \quad (11)$$

$$\mathbf{B}(\phi \wedge \psi) = \min\{\mathbf{B}(\phi), \mathbf{B}(\psi)\}, \quad (12)$$

$$\mathbf{B}(\phi \vee \psi) \geq \max\{\mathbf{B}(\phi), \mathbf{B}(\psi)\}. \quad (13)$$

C. Changing Beliefs

We can now investigate the degree of the agent's belief in terms of the labeling values of the arguments.

The following theorem states that a necessary and sufficient condition for formula ϕ to be believed to some extent is that, for all interpretation \mathcal{I} which does not satisfy ϕ , there exists an argument whose consequence is not satisfied by \mathcal{I} that is more accepted than every argument whose consequence is satisfied by \mathcal{I} .

Theorem 2: Let A , B , A_0 , and B_0 represent arguments, and let $\mu \in (0, 1]$ be a degree of belief. Then, for all $\phi \in \mathcal{L}$,

$$\begin{aligned} & \mathbf{B}(\phi) \geq \mu \\ \Leftrightarrow & \forall \mathcal{I} \not\models \phi \exists B_0 : \mathcal{I} \not\models \text{con}(B_0), \forall A : \mathcal{I} \models \text{con}(A), \\ & \alpha(B_0) - \alpha(A) \geq \mu. \end{aligned}$$

Proof:

$$\begin{aligned} & \mathbf{B}(\phi) \geq \mu \\ \Leftrightarrow & \forall \mathcal{I} \not\models \phi \quad \pi(\mathcal{I}) \leq 1 - \mu \text{ (Eq. 9)} \\ \Leftrightarrow & \forall \mathcal{I} \not\models \phi \\ & 1 + \max_{A:\mathcal{I} \models \text{con}(A)} \alpha(A) - \max_{B:\mathcal{I} \not\models \text{con}(B)} \alpha(B) \leq 1 - \mu \\ & \text{(Eq. 8)} \\ \Leftrightarrow & \forall \mathcal{I} \not\models \phi \quad \max_{B:\mathcal{I} \not\models \text{con}(B)} \alpha(B) - \max_{A:\mathcal{I} \models \text{con}(A)} \alpha(A) \geq \mu \\ \Leftrightarrow & \forall \mathcal{I} \not\models \phi \exists B_0 : \mathcal{I} \not\models \text{con}(B_0), \forall A : \mathcal{I} \models \text{con}(A), \\ & \alpha(B_0) - \alpha(A) \geq \mu. \end{aligned}$$

■

Corollary 1: A belief reinstatement of ϕ is total, i.e., the belief previously discarded is reintegrated with a degree $\mathbf{B}(\phi) = 1$, if and only if, for all interpretation \mathcal{I} which does not satisfy ϕ , there exists an argument whose consequence is not satisfied by \mathcal{I} that is fully accepted and every argument whose consequence is satisfied by \mathcal{I} is not accepted. Formally,

$$\begin{aligned} \mathbf{B}(\phi) &= 1 \\ \Leftrightarrow \forall \mathcal{I} \not\models \phi \exists B_0 : \mathcal{I} \not\models \text{con}(B_0), \forall A : \mathcal{I} \models \text{con}(A), \\ &\quad \alpha(B_0) - \alpha(A) = 1, \\ \Leftrightarrow \forall \mathcal{I} \not\models \phi \exists B_0 : \mathcal{I} \not\models \text{con}(B_0), \forall A : \mathcal{I} \models \text{con}(A), \\ &\quad \alpha(B_0) = 1, \alpha(A) = 0. \end{aligned}$$

There remains to be established the conditions under which belief in ϕ does not get reinstated at all.

The following theorem states that a formula ϕ is not (or no more) believed by an agent if and only if there exists an interpretation \mathcal{I}_0 which does not satisfy ϕ and is such that there exists an argument whose consequence is satisfied by \mathcal{I}_0 and is more accepted than all the arguments whose consequence is not satisfied by \mathcal{I}_0 .

Theorem 3: The necessary and sufficient condition for formula ϕ not to be believed may be stated as follows:

$$\begin{aligned} \mathbf{B}(\phi) &= 0 \\ \Leftrightarrow \exists \mathcal{I}_0 \not\models \phi, \exists A_0 : \mathcal{I}_0 \models \text{con}(A_0), \forall B : \mathcal{I}_0 \not\models \text{con}(B), \\ &\quad \alpha(A_0) \geq \alpha(B). \end{aligned}$$

Proof:

$$\begin{aligned} \mathbf{B}(\phi) &= 0 \\ \Leftrightarrow \exists \mathcal{I}_0 \not\models \phi : \pi(\mathcal{I}_0) = 1, \\ \Leftrightarrow \exists \mathcal{I}_0 \not\models \phi : \\ &\quad \min\{1, 1 + \max_{A:\mathcal{I}_0 \models \text{con}(A)} \alpha(A) - \max_{B:\mathcal{I}_0 \not\models \text{con}(B)} \alpha(B)\} = 1, \\ \Leftrightarrow \exists \mathcal{I}_0 \not\models \phi : \max_{A:\mathcal{I}_0 \models \text{con}(A)} \alpha(A) \geq \max_{B:\mathcal{I}_0 \not\models \text{con}(B)} \alpha(B). \end{aligned}$$

■

To sum up, if at some point belief in ϕ is lost due to the arrival of an acceptable argument A whose conclusion is satisfied by an interpretation $\mathcal{I}_0 \not\models \phi$, which causes the labeling (and therefore the belief set) to change so that $\mathbf{B}(\phi) = 0$, three things may happen when a new argument A' arrives:

- 1) $\mathcal{I}_0 \not\models \text{con}(A')$, and the new labeling is such that $\alpha(A') = 1$ and $\alpha(A) = 0$; in this case, ϕ is *fully reinstated*, i.e., $\mathbf{B}(\phi) = 1$;
- 2) $\mathcal{I}_0 \not\models \text{con}(A')$, and the new labeling is such that $\alpha(A') = \alpha(A) + \mu$, with $0 < \mu < 1$; in this case, ϕ is *partially reinstated*, i.e., $\mathbf{B}(\phi) = \mu$;
- 3) $\mathcal{I}_0 \models \text{con}(A')$, or $\mathcal{I}_0 \not\models \text{con}(A')$ but the new labeling is such that $\alpha(A') \leq \alpha(A)$; in both cases, there is *no reinstatement* of ϕ , and $\mathbf{B}(\phi) = 0$.

It should be noted, however, that, unlike in the crisp case, when (full or partial) reinstatement of a belief happens, the

reinstated belief may be stronger than, equal to, or weaker than the belief that had been lost. The actual outcome of reinstatement depends on how much the sources of the arguments responsible for the loss, in the first place, and the subsequent recovery of the belief are trusted by the agent.

This speaks in favor of the realism of our theory, which agrees with, and may be regarded as a possible explanation of, experimental results recently reported by Rahwan and colleagues [9]. In a number of psychological experiments, they found out that reinstatement does not always yield the full recovery of the attacked argument, which is to be expected if the source of the argument responsible for reinstatement is not fully trusted.

V. CONCLUSION

We have characterized the conditions under which the belief reinstatement is total, partial, or nonexistent in a fuzzy argumentation-based approach to possibilistic belief revision.

The proposed framework can be further improved following two directions: (i) specifying the trustworthiness degree by a cognitive model of trust and, (ii) embedding the proposed framework into an integrated one where also desires and, afterwards, goals are taken into account.

REFERENCES

- [1] C. da Costa Pereira, A. Tettamanzi, and S. Villata, "Changing one's mind: Erase or rewind? possibilistic belief revision with fuzzy argumentation based on trust," in *Procs of IJCAI*, 2011.
- [2] D. Dubois and H. Prade, "A synthetic view of belief revision with uncertain inputs in the framework of possibility theory," *Int. J. Approx. Reasoning*, vol. 17, no. 2-3, pp. 295–324, 1997.
- [3] C. da Costa Pereira and A. Tettamanzi, "An integrated possibilistic framework for goal generation in cognitive agents," in *Procs. of AAMAS*, 2010, pp. 1239–1246.
- [4] A. Dragoni and P. Giorgini, "Belief revision through the belief-function formalism in a multi-agent environment," in *Procs. of ATAL, LNCS 1193*, 1996, pp. 103–115.
- [5] J. Delgrande, D. Dubois, and J. Lang, "Iterated revision as prioritized merging," in *Procs. of KR*, 2006, pp. 210–220.
- [6] P. M. Dung, "On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games," *Artif. Intell.*, vol. 77, no. 2, pp. 321–358, 1995.
- [7] M. Caminada, "On the issue of reinstatement in argumentation," in *Procs. of JELIA, LNCS 4160*, 2006, pp. 111–123.
- [8] L. A. Zadeh, "Fuzzy sets as a basis for a theory of possibility," *Fuzzy Sets and Systems*, vol. 1, pp. 3–28, 1978.
- [9] I. Rahwan, M. I. Madakkatel, J.-F. Bonnefon, R. N. Awan, and S. Abdallah, "Behavioral experiments for assessing the abstract argumentation semantics of reinstatement," *Cognitive Science*, vol. 34, no. 8, pp. 1483–1502, 2010.