Toward a Revision for Multi-Context Systems

Aldo Franco Dragoni
Computer Science Institute
University of Ancona
dragon@inform.unian.it

Paolo Giorgini
Department of Information and
Communication Technology
University of Trento - Italy
pgiorgini@science.unitn.it

Abstract

The behavior of agents acting on an external world strongly depends on their ability of changing opinions and intentions with respect to changes in the environment. Although, multi-context systems have been successfully applied for defining agents mental states' architecture, no revision methods for multi-context systems have been proposed to deal with the dynamicity of the environment. In this paper we first introduce BDI agents modeled by multi-context systems, and then we propose an extension of a basic belief revision framework to multi-context systems.

1 Introduction

BDI agents (namely agents able to have Beliefs, Desires and Intentions) [10, 17, 21, 22, 24, 26] are supposed to have mental states, which contains explicit beliefs, desires and intentions about the environment, and about the other agents' beliefs, desires and intentions. Multi-context systems have been successfully used for defining such mental states'architecture [2, 4, 19]. These are defined in terms of different contexts (each of which contains a set of formulae closed under some inference rules) and a set of bridge rules for transferring information between contexts. Different contexts are used to represent different mental attitudes, such as beliefs, desires and intentions, and the interactions between these components can be specified by means of the bridge rules between the contexts.

The behavior of agents acting on an external world strongly depends on their ability of changing opinions and intentions with respect to changes in the environment. In the last two decades, Belief Revision as been defined as the process of rearranging a cognitive state in order to embody incoming information while preserving global consistency. Since the seminal, philosophical and influential works of Grdenfors et al. [1] ideas on belief revision have been progressively refined and ameliorated toward normative, effective and quasi-computable paradigms. Some of the main theoretical contributes have been:

- the distinction between the notion of "revision" and that of "updating";
- the notion of "epistemic entrenchment";
- the duality between syntactic and semantic approaches;
- the notion of "revision for finite bases";
- the notion of "revision as transmutation of partial epistemic rankings",

Side by side to this "symbolic" line of research, there has been also a "numerical" way to belief revision whose main contributes were the probabilistic, possibilistic, and the evidence-based approaches. In [5, 7] we proposed a belief revision framework that combines symbolic and numerical operations. It operates in ATMS-style (Assumption Truth Maintenance System) to treat the symbolic part of information, and in Dempster-Shafer style to treat their numerical part. One of main advantages of the proposed framework is that belief revision can be iterated, making it practical and usable for implementing agent systems.

The main goal of this paper is extending our belief revision approach to multi-context systems. The motivations behind are quite simple: if we model a mental state as a multi-context system and we want such a mental state being able to cope with the dynamicity of the external world, then we need to extend our single-theory (cognitive state) updating/revision mechanism to multi-context systems (mental states) revision methods.

The paper is structured as follows. Section 2 introduces multi-context systems and how to model agents' mental state using contexts. Section 3 presents the basic belief revision framework that we extend to multi-context systems in Section 4. Finally, Section 5 presents some conclusions and future work.

2 Modeling mental states with contexts

Agents are supposed to be characterized by *mental states*. We regard a mental state as a structure based on two primitive mental attitudes: *beliefs* and *desires*. Intuitively, intentions are what the agent desires to be true (or false) and also it believes it could be true (or false). The "could" means that the agent is able to

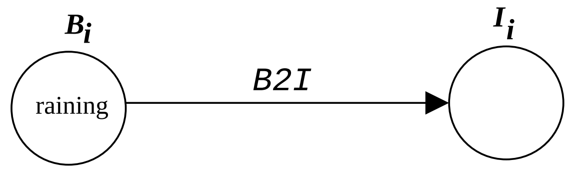

Figure 1: Contexts and bridge rules for agent i

act in order to change the external world and/or the other agents mental states to reach the desired state of affairs.

In this paper we introduce a very simple level of representation of mental states, just to give an example and justify our need of a multi-context systems revision method.

Following [4, 11, 12, 13], we use propositional contexts to formalize agents' mental states. To further simplify the matter, in this paper contexts are simply defined as propositional theories.

For any agent i, its sets of beliefs and intentions are represented by the contexts B_i and I_i, respectively. A formula ϕ in the context B_i (denoted by the pair $B_i : \phi$) represents the fact that i believes ϕ and, analogously, a formula ϕ in the context I_i ($I_i : \phi$) represents the fact that i has the intention to bring about ϕ.

In general, beliefs and intentions are not expressed in the same language. Although contexts support this possibility, for the sake of presentation we consider the simpler case in which for any agent i the languages for its beliefs and intentions coincide. We call this language L_i.

The beliefs and the intentions of an agent are not independent. The relation between the beliefs and the intentions of an agent can also be represented by bridge rules from the context of its beliefs to that of its intentions. For instance, the bridge rule:

$$\frac{B_i : raining}{I_i : bring\_umbrella}\ \mathcal{B}2\mathcal{I}$$

formalizes the fact that, if agent i believes that *it is raining*, then i intends *to bring an umbrella*. We indicate with $\mathcal{B}2\mathcal{I}$ the set of these bridge rules.

Figure 1 shows an example of contexts and bridge rules associated to the agent i (circles represent contexts and arrows represent bridge rules). In particular, i has a context for its beliefs (B_i), in which the formula *raining* is true, a context for its intentions (I_i), and the bridge rule $\mathcal{B}2\mathcal{I}$ connecting B_i and I_i.

Other contexts can be used for representing beliefs and intentions of other agents. For instance, in [8] we have introduced the idea of *image of mental states*, in which an agent i uses the contexts B_iB_j and B_iI_j to represent its beliefs regarding beliefs and intentions of another agent j. For the sake of simplicity, in this paper we consider only two contexts, B_i and I_i.

The logical systems that formalize the reasoning with a set of contexts connected by bridge rules are called *multi-context systems* [13].[1]

[1] In [13], multi-context systems are called multi-language systems to stress the fact that they

Definition 1 *A* multi-context system MC *is a pair* $\langle \mathsf{C}, \mathsf{BR} \rangle$*, where* C *is a set of contexts and* BR *a set of bridge rules. Any context in* C *is presented as an axiomatic propositional system* $\langle L, A \rangle$*, with* L *a propositional language and* $A \subseteq L$ *a set of proper axioms.*

The multi-context system associated to the structure of figure 1, is composed by the set of contexts $\mathsf{C} = \{B_h, I_h\}$, and the set of bridge rules $\mathsf{BR} = \{\mathcal{B}2\mathcal{I}\}$.

3 The Basic Belief Revision Method

In the following we present the sentence-based Belief Revision framework we have developed in [7] and that we want to extend to the multi-context case. For an agent i modeled with contexts, as presented in the previous section, the belief revision framework is applied locally to the context representing the agent's beliefs (B_i). Namely, the revision process is applied only to the context B_i without considering the effects produced by the bridge rules into the other contexts.

Our approach conceives two knowledge repositories:

1. the *knowledge background* B_i^*, which is the set of all the propositional sentences available to the reasoning agent i (as assumptions). Since, it can be inconsistent, it cannot be used as a whole to support reasoning and decision processes;

2. the *knowledge base* $B_i \subseteq B_i^*$, which is the maximally consistent, currently preferred piece of knowledge that should be used for reasoning and decision supporting.

The reason why we pick *maximally* consistent subsets of B_i^* is that this is the perfect implementation of the "Principle of Recoverability" [7] : everything were known in the past should be believed again now (or still believed) if nothing prevents that.

Following the idea of numerical approaches to belief revision, our approach associates to each belief a weight of credibility. Redefinition of these weights in the light of the incoming information is a crucial part of belief revision. Computationally, when an agent i acquires a new information p our belief revision mechanism follows four basic steps (Figure 2) [5, 6]:

S1 Detection of the minimally inconsistent subsets of $B_i^* \cup \{p\}$ (*nogoods*)

S2 Generation of the maximally consistent subsets of $B_i^* \cup \{p\}$ (*goods*)

S3 Revision of the credibility weights of the sentences in $B_i^* \cup \{p\}$

S4 Choice of a preferred *good* as the new revised base B_i'

allow for multiple distinct languages.

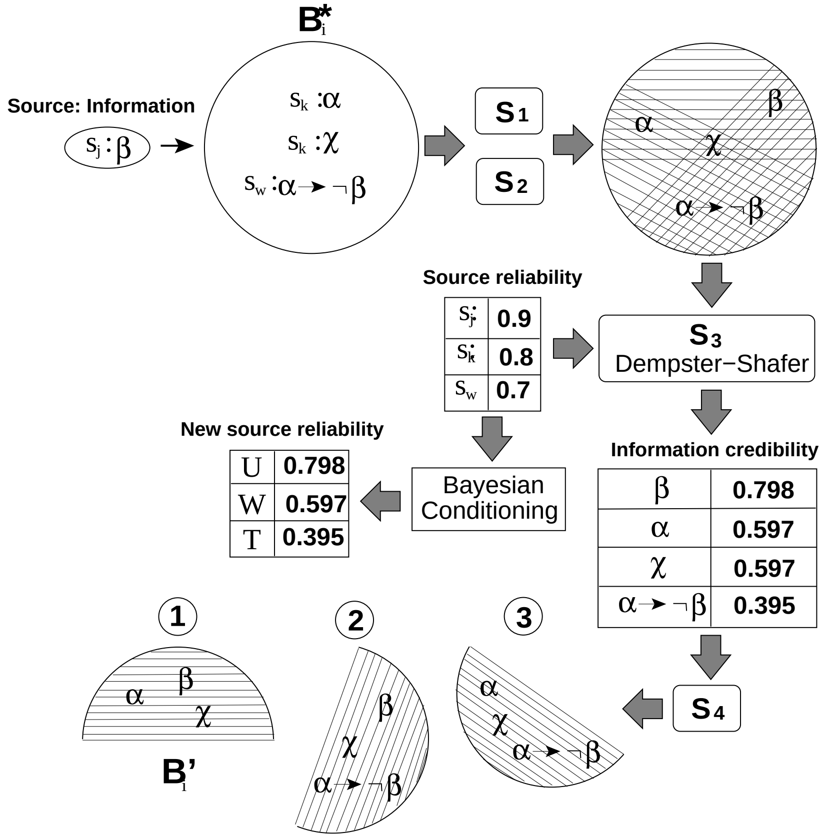

Figure 2: Dempster's and Bayes rules at work

The incoming information p, with its weight of evidence, is evaluated not just within the current base B_i, but within the overall knowledge background B_i^*. Doing so, the degrees of credibility of the sentences in $B_i^* \cup \{p\}$ are reviewed on a broader and less prejudicial basis (S3). The main advantage is that we can rescue sentences from B_i^* by virtue of the maximal consistency of B_i'. If we would revise only B_i by p, we could not recover information from B_i^*.

S1 and S2 deal with consistency and derivation, and act on the symbolic part of the information. Operations are in ATMS style; to find out nogoods and goods, we adopt (and adapt) the most efficient set-covering algorithm that we are aware of [23]. Notwithstanding this, even in the propositional case, determining all the minimal inconsistencies can be very hard. However, such condition can be relaxed (the consequence is that some of the goods are not really consistent) and in practical applications dealing with commonsense knowledge (for instance [9]), such minimal inconsistencies could be provided interactively by the external user.

S3 and S4 deal with uncertainty and work with the numerical weight of the information. Both contribute to the choice of the revised knowledge base so their reasonableness should be evaluated as a couple. Numerical formalisms are able to perform both of them since the credibility of a single sentence p is determined in the same way as the credibility of a set of sentences B_i by the weights attached to their models (or possible worlds) $[p]$ and $[B_i]$, respectively. Flexibility is an advantage in separating the two steps; for instance, depending on the characteristics of the knowledge domain under consideration and the kind of task and/or decision that should be taken on the basis of the revision outcome, the selection function could consider also one (or a combination) of the methods described in [3].

Probabilistic methods with uncertain inputs seem inadequate for the strong dependence that they impose on the credibility of a sentence and that of its negation. We see that the belief-function formalism, in the special guise in which Shafer and Srivastava apply it to auditing [25], could work well because it treats all the pieces of information as they had been provided at the same time.

The method has the following I/O:

INPUT: list of pairs <source,piece of information>
list of pairs <source,reliability>
OUTPUT: list of pairs <piece of information,credibility>
list of pairs <source,reliability>

Let S=$\{s_1, ..., s_n\}$ be the set of the sources, and let b_j^* be the subset of B_i^* that agent i has received from s_j. Each source s_j is associated with a *reliability* $R(s_j)$, that is regarded as the *probability* that the source is faithful. The main idea with this multi-source version of the belief function framework is that a reliable source cannot give false information, while an unreliable source can give correct information; the hypothesis that s_j is reliable is compatible only with the models of b_j^*, while the hypothesis that s_j is unreliable is compatible with the overall set of models Ω of L (in Depster-Sahfer theory, it is called *frame of discernment* [25]). Each source s_j is an evidence for B_i^* and generates the following *basic probability assignment* (*bpa*) $m_j(\cdot)$ on 2^Ω:

$$m_j(X) = \begin{cases} R(s_j) & \text{if } X = [b_j^*] \\ 1 - R(s_j) & \text{if } X = \Omega \\ 0 & \text{otherwise} \end{cases}$$

All these *bpa*s will be then combined through the Dempster Rule of Combination. From the combined *bpa* $m(\cdot)$, the credibility of a sentence p of L is given, as usual, by:

$$Bel(p) = \sum_{X \subseteq [p]} m(X)$$

From this mechanism we obtained an easy way to calculate the new reliabilities of the sources. Let Φ be an element of 2^S. If the sources are independent, the reliability of Φ is

$$R(\Phi) = \prod_{S \in \Phi} R(s) \cdot \prod_{S \in \Phi} (1 - R(s))$$

It holds that

$$\sum_{\Phi \in 2^S} R(\Phi) = 1$$

It maybe that some source fall in contradiction, so that some elements of 2^S are impossible. The remaining elements are subjected to bayesian conditioning so that their reliabilities sum up again to 1. The revised reliability $R'(s_j)$ of a source s_j is the sum of the new reliabilities of the surviving elements of 2^S that contain s_j. If a source has been involved in some contradictions, then $R'(s_j) \leq R(s_j)$, otherwise $R'(s_j) = R(s_j)$.

The main problem with the belief function formalism is the computational complexity of Dempster's Rule of combination. The straight-forward application of the rule is exponential in the frame of discernment (number of propositional letters of L, that is smaller than the number of information items in B_i^*) and the number of evidences. However, much effort has been spent in reducing the complexity of the rule. Such methods range from "efficient implementations" [14] to "qualitative approaches" [18] through "approximate techniques" [16].

Among the "quantitative" methods to perform S4, we have chosen to order the goods according to the average credibility of their elements. A main difference with respect to other methods like *best-out method*, *inclusion-based method*, and *lexicographic method*, is that the preferred good(s) may no longer necessarily contain the most credible piece(s) of information (see [7] for a more detailed discussion).

A final step in our revision mechanism is the selection of the derived sentences which are still derivable from B_i' since the assumptions on which they rely are all contained in B_i''. Theoretically, it simply consists in applying classical entailment on the preferred good to deduce plausible conclusion from it. We adopted an ATMS and we stored each sentence derived by the Theorem Prover with an *origin set* [15], i.e., a set of basic assumptions which are all *necessary* to derive it. Practically, this step (not represented in figure 2) consists in selecting from the derived sentences, all those whose origin set is subset of the preferred good. We could relax the definition of origin set to that of a set of basic assumptions used to derive the sentence. This is easier to compute and does not have harmful consequences; the worst it can happen is that, being this relaxed origin set a superset of the real one, it is not certain that it will be a subset of the preferred good as the real one is, and so some derived logical consequences of the preferred good may be not recognized (at first).

Besides recoverability, this computational model for belief revision overcomes various limitations of other classic approaches. In particular, the revision can be iterated, it is more flexible, and the splitting between the symbolic treatment of the inconsistencies and the numerical revision of the credibility weights, provides a clear understanding of what is going on and lucid explanations for the choices ([7] for more details).

4 Multi-Context Revision

In this section we present our approach to the multi-context revision. The idea is to extend the belief revision model presented in the previous section to the case of a multi-context system. We want this approach to be sentence-based and to respect the "recoverability principle".

The correspondence between Mono-Contextual and Multi-Contextual revision is sketched in Table 1.

Mono-context Revision	**Multi-context Revision**
Knowledge Background (B_i^* inconsistent)	*Multi-Context Background* (MCB inconsistent)
Theory to be revised.	This is the global multi-context system that collects all the contexts. Contexts are fixed and predefined. Bridge rules among contexts and assumptions (sentences) inside them are introduced incrementally. In compliance with the "recoverability principle", nothing will ever be removed from MCB neither assumptions nor bridge rules.
Good	*Scenery*
Maximally consistent subset of B_i^*.	This is a subsystem of MCB made of all its contexts but, eventually, without some assumptions and/or some bridge rules, in such a way that it results maximally consistent. This means that it is not possible to re-introduce any assumption and/or bridge rule without generating a contradiction in some contexts.

Table 1: Correspondence between mono-context and multi-context revision

Let us try to be a bit more precise. In the multi-context system we consider fixed the set of contexts, lets call L the set of labels associated to each context (for instance in the example of Figure 1, $\mathsf{L}= \{B_i, I_i\}$). The sentences in these contexts called "assumptions", and the bridge rules between them are collected as the time goes on. The assumptions play, as usual, the same role that proper axioms play in a logical theory. Let us call KB the set of all these assumptions (a difference w.r.t. the Mono-Contextual case is that here assumptions are *labeled* sentences), and let us call BR the set of all the Bridge Rules introduced so far. We define Multi-Context Background the tuple:

$$\mathsf{MCB} = \langle \mathsf{L}, \mathsf{KB}, \mathsf{BR} \rangle$$

Of course, MCB could fall in contradiction (when at least a context becomes inconsistent). When this happens, MCB will be "partitioned" into a finite set of

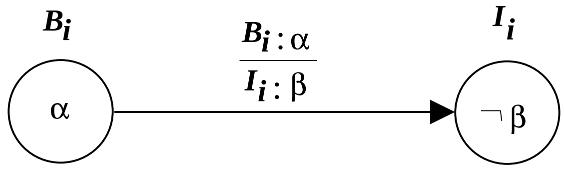

Figure 3: MCB for the agent i

subsystems, called "sceneries".

A scenery is a couple $\langle B, BR \rangle$ where $B \subseteq \mathsf{KB}$ and $BR \subseteq \mathsf{BR}$ such that:

- $\langle \mathsf{L}, B, BR \rangle$ is consistent;
- $\forall B' \ (B \subset B' \subseteq \mathsf{KB}),\ \langle \mathsf{L}, B', BR \rangle$ is inconsistent;
- $\forall BR' \ (BR \subset BR' \subseteq \mathsf{BR}),\ \langle \mathsf{L}, B, BR' \rangle$ is inconsistent too.

Lets consider the example in Figure 3, in which a multi-context background for agent i presents a contradiction in the context I_i (both β and $\neg\beta$ are true).

To solve the contradiction, the three sceneries in Figure 4 are generated. In scenery 1, $\neg\beta$ is removed from I_i; in scenery 2, α is removed from B_i, and finally in scenery 3, the bridge rule is removed.

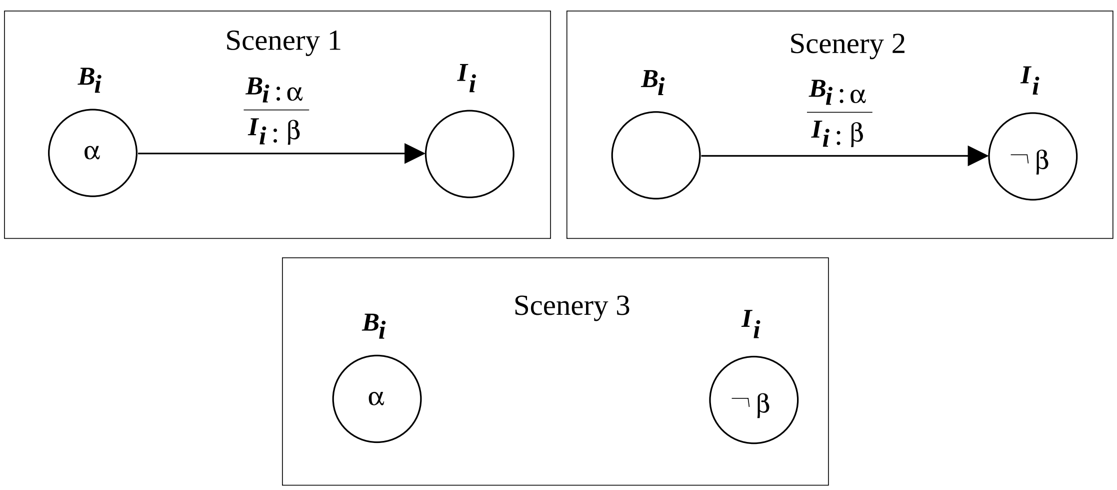

Figure 4: Sceneries for the example in Figure 3

Following the trivial example in Section 2 (the umbrella example), let us suppose that a new intention is introduced into the agent's mental state, i.e., that of not to bring an umbrella. We should introduce the assumption: $I_i : \neg bring\_umbrella$ in KB. This renders MCB inconsistent. The revision method yields the following three sceneries:

1. $\langle \{B_i : raining\}, \frac{B_i : raining}{I_i : bring\_umbrella} \mathcal{B}2\mathcal{I} \rangle$

2. $\langle\{I_i : \neg bring\_umbrella\}, \frac{B_i : raining}{I_i : bring\_umbrella}\mathcal{B}2\mathcal{I}\rangle$

3. $\langle\{I_i : \neg bring\_umbrella, B_i : raining\}\rangle$.

It is easy to see that any other subsystem of MCB is not maximally consistent. In Figure 4 the pseudo-algorithm for generating all the sceneries.

Main:

1. collect all the context labels in the set L;
2. collect all the bridge rules in the set BR;
3. initialize KB $= \emptyset$;
4. initialize $Sceneries = \{\mathsf{L}\}$;
5. repeat
6. read($NewInformation$);
7. IF $NewInformation = (Assumption, Context)$
   THEN $revise(Assumption, Context)$
   ELSE $revise(BridgeRule)$
8. endrepeat
9. end

PROCEDURE $revise(Assumption, Context)$

1. while $Sceneries \neq \emptyset$ do begin
   (a) POP($S, Sceneries$);
   (b) IF $derivable((\neg Assumption, Context), \mathsf{BR}, S) = true$
      THEN
         i. $generate\_sceneries(\mathsf{BR}, \mathsf{KB} \cup \{(Assumption, Context)\}, New\_Sceneries)$;
         ii. $Sceneries := Sceneries \cup NewSceneries$;
      ELSE
         i. $Context := Context \cup \{Assumption\}$;
         ii. $update(S, Context)$;
2. endwhile
3. $eliminate\_subscenaries(Sceneries)$
4. end

Figure 5: The pseudo-algorithm for generating all the sceneries

5 Conclusion

This paper is a preliminary step toward extending a classic sentence-based revision mechanism from single theories to multi-context systems. After introducing the motivations behind that, the paper presents the main idea of a Multi-Context

Revision method. It also contains a preliminary version of a main algorithm to perform the task (in a decidable multi-context system). Most of the work has still do be done. Here is a list of theoretical and practical task to do.

1. Prove that at least one scenery exists
2. Prove that the number of generated sceneries is finite
3. Prove that the process is deterministic (the set of generated sceneries is unique)
4. Find an algorithm for $revise(BridgeRule)$
5. Find an algorithm for $derivable((Assumption, Context), \mathsf{BR}, S)$
6. Find an algorithm for $generate\_sceneries(\mathsf{BR}, \mathsf{KB} \cup \{(A, C)\}, NS)$
7. Find an algorithm for $eliminate\_subsceneries(Sceneries)$
8. Prove that the overall algorithm is complete
9. Prove that the overall algorithm is sound
10. Implement the overall algorithm

References

[1] C.E. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change: Partial meet contraction and revision functions. *The Journal of Simbolic Logic*, 50:510–530, 1985.

[2] M. Benerecetti, F. Giunchiglia, and L. Serafini. Model Checking Multiagent Systems. *Journal of Logic and Computation, Special Issue on Computational & Logical Aspects of Multi-Agent Systems*, 8(3):401–423, 1998.

[3] S. Benferhat, C. Cayrol, D. Dubois D., J. Lang, and H. Prade. Inconsistency management and prioritized syntax-based entailment. In *Proc. of the 13th Inter. Joint Conf. on Artificial Intelligence*, pages 640–645, 1993.

[4] A. Cimatti and L. Serafini. Multi-Agent Reasoning with Belief Contexts II: Elaboration Tolerance. In *Proceedings of the 1st International Conference on Multi-Agent Systems (ICMAS-95)*, pages 57–64, 1996.

[5] A.F. Dragoni and P. Giorgini. Belief revision through the belief function formalism in a multi-agent environment. In Wooldridge M., Jennings N.R., and Muller J., editors, *Intelligent Agents III*, number 1193 in LNCS. Springer-Verlag, 1997.

[6] A.F. Dragoni and P. Giorgini. Distributed knowledge revision-integration. In *Proc. of the Sixth ACM International Conference on Information Technlogy and Management*. ACM Press, 1997.

[7] A.F. Dragoni and P. Giorgini. Revising beliefs received from multiple source. In M A Williams and H Rott, editors, *Frontiers of Belief Revision*, Applied Logic. Kluwer, 2001.

[8] A.F. Dragoni, P. Giorgini, and L. Serafini. Mental States Recognition from Communication. *Journal Logic and Computation*, (to appear).

[9] A.F. Dragoni and M. Di Manzo. Supporting complex inquiries. *International Journal of Intelligent Systems*, 10:959–986, 1995.

[10] M. Georgeff. Communication and interaction in multiagent planning. In *Proceedings of the 3th National Conference on Artificial Intelligence*, pages 125–129, 1983.

[11] F. Giunchiglia. Contextual reasoning. *Epistemologia, special issue on I Linguaggi e le Macchine*, XVI:345–364, 1993.

[12] F. Giunchiglia and C. Ghidini. Local Models Semantics, or Contextual Reasoning = Locality + Compatibility. In *Proceedings of the Sixth International Conference on Principles of Knowledge Representation and Reasoning (KR'98)*, pages 282–289. Morgan Kaufmann, 1998.

[13] F. Giunchiglia and L. Serafini. Multilanguage hierarchical logics (or: how we can do without modal logics). *Artificial Intelligence*, 65:29–70, 1994.

[14] R. Kennes. Computational aspects of the möbius transform of a graph. *IEEE Transactions in Systems, Man and Cybernetics*, 22:201–223, 1992.

[15] J.P. Martins and S.C. Shapiro. A model for belief revision. *Artificial Intelligence*, 35:25–97, 1988.

[16] S. Moral and N. Wilson. Importance sampling monte-carlo algorithms for calculation of dempster-shafer belief. In *Proc. of IPMU'96*, Granada, 1996.

[17] J.S. Rosenschein M.R. Gensereth, M.L. Ginsberg. Cooperation without communication. In *AAAI 86*, pages 51–57, 1986.

[18] S. Parsons. Some qualitative approaches to applying the dempster-shafer theory. *Information and Decision Technologies*, 19:321–337, 1994.

[19] S. Parsons, C. Sierra, and N. R. Jennings. Agents that reason and negotiate by arguing. *Journal of Logic and Computation*, 3(8):261–292, 1998.

[20] D. Prawitz. *Natural Deduction - A proof theoretical study.* Almquist and Wiksell, Stockholm, 1965.

[21] A.S. Rao and M. Georgeff. BDI agents: from theory to practice. In *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95)*, pages 312–319, S. Francisco, CA, 1995.

[22] A.S. Rao, M. Georgeff, and E.A. Sonenberg. Social plans: A preliminary report. In E. Werner and Y. Demazeau, editors, *Decentralized AI - Proceedings of the Third European Workshop on Modeling Autonomous Agents in a Multi-Agent World (MAAMAW-91)*, pages 57–76, Amsterdam, The Netherlands, 1992. Elsevier Science Publishers B.V.

[23] R. Reiter. A theory of diagnosis from first principles. *Artificial Intelligence*, 53, 1987.

[24] J.S. Rosenschein and M.R. Genesereth. Communication and cooperation. *Stanford Heuristic Programming Rep*, 1984.

[25] G. Shafer and R. Srivastava. The bayesian and belief-function formalisms a general perspective for auditing. In G. Shafer and J. Pearl, editors, *Readings in Uncertain Reasoning*. Morgan Kaufmann, 1990.

[26] E. Werner. Toward a theory of communication and cooperation for multiagent planning. In *Proceedings of the Second Conference on Theoretical Aspects of Reasoning About Knowledge*, Los Altos, CA, 1988. Morgan Kaufmann Publisher.