
Mental States Recognition from Communication

ALDO FRANCO DRAGONI, *Istituto di Informatica, University of Ancona, Via Brezze Bianche, 60131, Ancona, Italy.*

E-mail: dragon@inform.unian.it

PAOLO GIORGINI, *Faculty of Science, University of Trento, Via Sommarive, 14, 38050 Povo - Trento, Italy.*

E-mail: pgiorgini@science.unitn.it

LUCIANO SERAFINI, *Luciano Serafini, ITC-IRST, Povo, 38050, Trento, Italy.*

E-mail: serafini@irst.itc.it

Abstract

In order to perform effective communication, agents must be able to foresee the effects of their utterances on the addressee's mental state. In this paper we study the consequences of an utterance on the mental state of a hearer. Given an agent communication language with a STRIPS-like semantics, we propose a set of criteria that allow the binding of the speaker's mental state to its uttering of a certain sentence. On the basis of these criteria, we give an abductive procedure that the hearer can adopt to partially recognize the speaker's mental state that led to a specific utterance.

Keywords: Abduction, belief revision, speech acts.

1 Introduction

In multi-agent systems, if the agents are not designed with embedded pre-compiled knowledge about the beliefs, intentions, abilities and perspective of other agents, they need to exchange information in order to cooperate and coordinate their activities. However, in real application domains, communication might be a limited resource (limited bandwidth, low signal/noise ratio etc.). In such cases, it is very important that, when deciding whether to send a message, agents consider their expected benefits vs. the costs of communication [17].

In order to evaluate the utility of communication, agents must be able to foresee the effects of their utterances on the addressee's mental state. To this end, a speaker can try to recognize and model the hearer's mental state and then predict the impact of its messages on the basis of this model. Mental states recognition becomes then very important for agents to perform effective communication. This task can be accomplished by both the observation of the other agents' behaviour [30] and the content of their utterances [6]. Assuming that the communication is the main road for exchanging portions of mental states, in this paper we focus on how to recognize the speaker's mental state on the basis of its utterances.

BDI agents (namely agents able to have beliefs, desires and intentions) [10, 22, 28, 29, 31, 33] are supposed to have a mental state, which contains beliefs, desires and intentions about

the environment, and about the other agents' beliefs, desires and intentions. The behaviour of an agent strongly depends on its mental state and communication is generally used to affect the behaviour of other agents. Communication is supposed to be intentional, i.e. activated by the speaker's reasoning about its own beliefs, desires and intentions. In other words, it is generally possible to regard utterances as the consequence of the speaker's being in a particular mental state, that provokes the desire to influence the hearer's mental state. This position is very general and is independent of the particular class of speech acts (assertive, commissive, directive, declarative or expressive) under consideration. If utterances are effects of a forward reasoning based on mental conditions, it seems natural to suppose that the hearer can use the received communication as base of a backward reasoning for recognizing the hypothetical speaker's mental state that originated the communication. Abductive reasoning can be adopted for this end. In particular, the hearer can use abduction as backward reasoning from the kind and the content of the received communication to the speaker's mental state.

The main goal of this paper is to provide some methods that a hearer can adopt in order to recognize and update the speaker's mental state. To do this we propose:

1. a formal representation of mental states based on the theory of contexts [4];
2. a correlation between the mental state of an agent and its utterances, based on the plan-based theory of speech acts [5];
3. a formal framework that generalizes the ideas of abductive reasoning to multi-context systems;
4. a formal characterization of the operations for updating the hearer's mental state.

The novelty of the paper is the fact that, not only do we devise an abductive theory for revising the mental state of an agent, but we also relate this theory to the semantics of agent communication languages. Differently from the major agent communication languages in which it is not specified how their semantics must be used by the agents, we propose a computationally feasible way to treat a set of communicative acts with semantics expressed in terms of preconditions and main effects.

The paper is structured as follows. In Section 2 we present the context framework, which we use to formalize agents' mental states. Section 3 describes a simple scenario used as an explanatory example throughout the paper. Section 4 illustrates how to exploit the plan-based theory of speech acts to correlate the speaker's mental state to its utterances. In Section 5 we recall Konolige's definition of causal theories and abduction, and we extend it to multi-context systems. We define three basic operations on mental states: abductive expansion, abductive contraction and abductive revision. Section 6 presents some abductive methods which can be used to update the hearer's image of the speaker's mental state. Finally, Section 7 discusses the related work, and Section 8 presents some conclusions and future work.

2 Mental states

Agents are supposed to be characterized by *mental states*. Mental states change as agents interact with one other and each change is clocked by the execution of an action. We limit our attention to the understanding of what changes between two successive snapshots of a mental state, before and after the execution of a 'Speech Act'. Previous work [6] proposes abductive methods which allow one to hypothesize the mental conditions behind the uttering of a speech act. The agents' mental state is formalized by modal logics, but unfortunately, the notion of abduction is not extended to modal logic. Here, we circumvent that incongruity

by exploiting certain results in modal logics and in propositional hierarchical logic contained in [15].

2.1 Modelling agents with contexts

We regard a mental state as a structure based on two primitive mental attitudes: *beliefs* and *desires*. Intuitively, intentions represent 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 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 do not distinguish between desires and intentions since we suppose that whenever an agent desires to change the hearer’s mental state and it believes that this is possible by means of a communicative act, then it intends to do it; and vice versa, if the agent intends to change the hearer’s mental state, then it also desires to do it.

Following [4, 13, 14, 15], we use propositional contexts to formalize agents’ mental states. A context is defined as a set of formulae closed under a set of inference rules (a theory).

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 . L_i contains formulae to represent the environment and formulae to express the fact that an agent has a certain belief or intention. In particular, for any agent j , L_i contains the following atomic propositions:

$B_j\phi$, meaning that agent j believes ϕ ;

$I_j\phi$, meaning that agent j intends to bring about ϕ .

Formulae of the form $B_j\phi$ and $I_j\phi$ are called BDI atoms [3].

Reasoning capabilities of an agent i on its beliefs and intentions, are represented in the contexts B_i and I_i by two sets of inference rules. Examples of reasoning capabilities could be any set of logical inference rules. Reasoning capabilities are supposed to be general purpose reasoning machineries, which do not contain special inference rules for BDI atoms. For the sake of this paper we suppose that any inference machinery is the set of rules for propositional logic.

In B_i and I_i , BDI atoms are considered as any other atomic formula. This implies, for instance, that $B_j\phi$ and $B_j(\phi \vee \psi)$ are completely independent beliefs of i . On the other hand, if i ascribes to the agent j enough reasoning capabilities, then either i believes that j believes both ϕ and $\phi \vee \psi$, or i believes that j believes neither ϕ nor $\phi \vee \psi$. The relation among BDI atoms in B_i and I_i , therefore, depends on the reasoning capabilities that i ascribes to j . The beliefs, the intentions, and the reasoning capabilities that i ascribes to another agent j are explicitly modelled by means of a mental state, called i ’s *image of j ’s mental state*. In particular, i ’s beliefs about the beliefs and the intentions of j , are represented by the contexts B_iB_j and B_iI_j , respectively. The same representation is used to formalize i ’s intentions regarding j ’s beliefs and intentions, that is the contexts I_iB_j and I_iI_j . Analogously, i ’s beliefs about j ’s beliefs about another agent k are formalized by the pair of contexts $B_iB_jB_k$ and $B_iB_jI_k$. This iteration can go on infinitely, but, in the cases of artificial agents with finite

resources, it is more adequate to consider a finite amount of iterations. On the other hand, we do not put any upper-bound on the limit of nested beliefs and intentions. For the sake of the explanation we consider only three levels of nesting.

The intuitive interpretation of a formula depends on the context. For instance, as already said, the formula ϕ in the context B_i , denoted by $B_i : \phi$, expresses the fact that agent i believes ϕ . The same formula in the context $B_i B_j I_i$, denoted by $B_i B_j I_i : \phi$, expresses the fact that i believes that j believes that i intends ϕ . On the other hand, different formulas in different contexts can represent the same fact. For instance, the formulas ϕ in the context $B_i I_j$ and the formula $I_j \phi$ in the context B_i have the same meaning. The effect of this ‘meaning overlapping’ is that, contexts cannot be considered as isolated theories, and that the set of theorems of a context might affect the set of theorems in another context. The interaction among contexts is formalized by *bridge rules*. In particular, we use the following bridge rules between contexts in a mental state and contexts in images of mental state:

$$\frac{\alpha : B_i \phi}{\alpha B_i : \phi} \mathcal{R}_{dn.B} \quad \frac{\alpha B_i : \phi}{\alpha : B_i \phi} \mathcal{R}_{up.B} \quad \frac{\alpha : I_i \phi}{\alpha I_i : \phi} \mathcal{R}_{dn.I} \quad \frac{\alpha I_i : \phi}{\alpha : I_i \phi} \mathcal{R}_{up.I}$$

$\mathcal{R}_{dn.B}$ and $\mathcal{R}_{dn.I}$ are called *reflection down*, whereas $\mathcal{R}_{up.B}$ and $\mathcal{R}_{up.I}$ *reflection up*. Reflection up and reflection down are often used in combination. For instance, reflection down allows an agent i to convert a formula $B_j \phi$ into a simpler format (by eliminating B_j) in order to perform reasoning about ϕ in its image of j ’s mental state. Reflection up is used by i to lift up, in its mental state, the result of such a reasoning. This reasoning pattern allows i to prove relations among BDI atoms. For instance, to prove $B_j \phi \supset B_j(\phi \vee \psi)$ in the context of its beliefs, i can perform the following deduction:

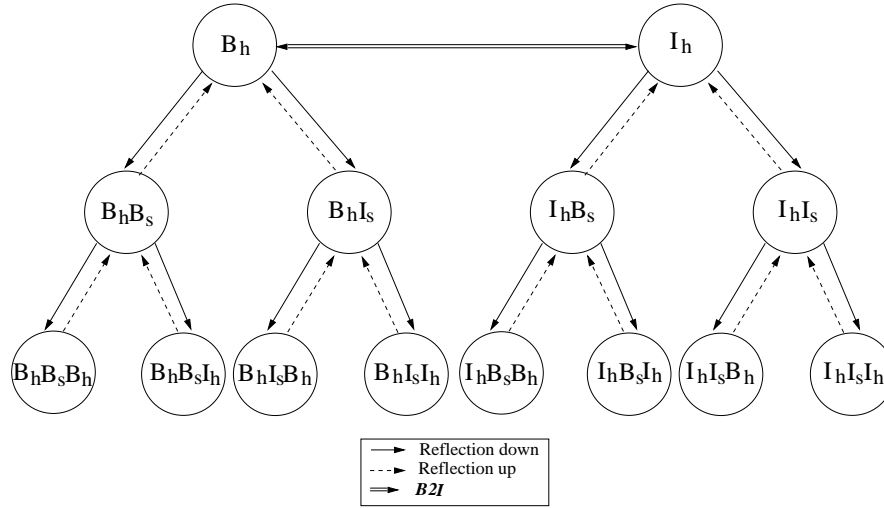
$$\frac{\frac{\frac{B_i : B_j \phi}{B_i B_j : \phi} \mathcal{R}_{dn.B} \quad \vee I}{B_i B_j : \phi \vee \psi} \mathcal{R}_{up.B}}{B_i : B_j(\phi \vee \psi)} \supset I$$

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:

$$\frac{B_i : \textit{raining}}{I_i : \textit{bring_umbrella}} \mathcal{B2I}$$

formalizes the fact that, if agent i believes that *it is raining*, then i intends to *bring an umbrella*. We indicate with $\mathcal{B2I}$ the set of these bridge rules.

Since we are interested in formalizing the effects of an utterance performed by a speaker on the beliefs and intentions of a hearer, we consider the two agents s and h denoting the speaker and hearer, respectively. Furthermore, we focus only on the effects of the utterance on the hearer’s mental state. We therefore consider only the contexts for the hearer’s beliefs and intentions (namely B_h and I_h), the contexts for the hearer’s beliefs and intentions regarding the speaker’s beliefs and intentions (namely $B_h B_s$, $B_h I_s$, $I_h B_s$, and $I_h I_s$), and the contexts for the hearer’s beliefs and intentions regarding the speaker’s beliefs and intentions regarding the hearer’s beliefs and intentions. Of course, this nesting could be extended indefinitely (for


 FIGURE 1. Contexts for agent h

more details see [11, 12, 16]), but three levels (as depicted in Figure 1, where circles represent contexts and arrows represent bridge rules) are sufficient to illustrate the abductive methods for the inference of mental states from communicative acts.

2.2 Mental states

The logical systems that formalize the reasoning with a set of contexts connected by bridge rules are called *multi-context systems* [15].¹

DEFINITION 2.1

A *multi-context system* MC is a pair $\langle C, BR \rangle$, where C is a set of contexts and BR a set of bridge rules. Any context in C is presented as an axiomatic formal system $\langle L, A, R \rangle$, with L a logical language, $A \subseteq L$ a set of axioms, and R a set of inference rules.

Derivability in a MC , in symbols \vdash_{MC} , is defined in terms of Natural Deduction [27]. Deductions in MC are trees of wffs starting from a finite number of formulae (either axioms or assumptions), possibly belonging to distinct contexts, and applying a finite number of inferences and bridge rules of BR . A wff $\alpha : \phi$ is derivable from the set of axioms AX in MC (in symbols, $AX \vdash_{MC} \alpha : \phi$), if there is a deduction that ends with $\alpha : \phi$ and whose axioms are in AX . For a detailed description on the proof theory of MC we refer the reader to [15].

DEFINITION 2.2

Let MC a multi-context system $\langle C, BR \rangle$. For each context $\alpha \in C$ and a set of axioms AX , α^* is the set of all formulae of α derivable from the set of axioms AX in MC . In symbols, $\alpha^* = \{\phi \mid AX \vdash_{MC} \alpha : \phi\}$.

In our special case, the multi-context system associated to the structure of figure 1, is composed by the set of contexts $C = \{B_h, I_h, B_h B_s, B_h I_s, I_h B_s, I_h I_s, B_h B_s B_h, B_h B_s I_h,$

¹In [15], multi-context systems are called multi-language systems to stress the fact that they allow for multiple distinct languages.




FIGURE 2. The MC of the working example

$B_h I_s B_h, B_h I_s I_h, I_h B_s B_h, I_h I_s I_h$ }, and the set of bridge rules $BR = \{\mathcal{R}_{up.B}, \mathcal{R}_{up.I}, \mathcal{R}_{dn.B}, \mathcal{R}_{dn.I}, \mathcal{B2I}\}$. Furthermore, we suppose that for each context, the language is a propositional language, as described in previous section, the set of axioms is empty, and the inference rules are the natural deduction inference rules for propositional logic.

To describe the effects of communication, we need to represent the hearer's mental state at different points in time. For each point in time, we consider a set of axioms AX . Intuitively, AX represents the basic beliefs, intentions, and nested beliefs and intentions of the hearer at that point. All the other formulae can be deduced in MC starting from AX . For instance, the beliefs of the hearer are the formulae in B_h^* , its intentions are the formulae in I_h^* , its beliefs about the speaker's beliefs are the formulae in $B_h B_s^*$, and so on.

The effects of the receipt of a message from the speaker is represented in the hearer by a change of AX into a new set of axioms AX' . Extending Definition 2.2, for each context α of MC, we define the set $\alpha'^* = \{\phi \mid AX' \vdash_{MC} \alpha : \phi\}$.

DEFINITION 2.3

The *mental state* of the agent h is the pair of the sets containing h 's beliefs and h 's intentions derivable in MC. In symbols, $ms(h) = \langle B_h^*, I_h^* \rangle$.

DEFINITION 2.4

h 's *image of s ' mental state* is the pair of sets containing h 's beliefs on s ' beliefs and h 's beliefs on s ' intentions, derivable in MC. In symbols, $ms(h, s) = \langle B_h B_s^*, B_h I_s^* \rangle$.

DEFINITION 2.5

h 's *image of s ' image of h 's mental state* is the pair of sets containing h 's beliefs on s ' beliefs on h 's beliefs and h 's beliefs on s ' beliefs on h 's intentions, derivable in MC. In symbols, $ms(h, s, h) = \langle B_h B_s B_h^*, B_h B_s I_h^* \rangle$.

3 Working example

Let consider the multi-context system MC represented in Figure 2 (the labels indicate both the formulae that belong to the contexts and the bridge rules) and characterized by the following $\mathcal{B2I}$ bridge rule and set of axioms AX .

$$\begin{array}{c}
\frac{B_h : B_s \neg \text{conditioning_on}}{B_h B_s : \neg \text{conditioning_on}} \mathcal{R}_{dn.B} \quad B_h B_s : (3.5) \\
\hline
\frac{B_h B_s : \text{temp_higher_}20^\circ}{B_h : B_s \text{temp_higher_}20^\circ} \mathcal{R}_{up.B} \quad B_h : (3.2) \\
\hline
\frac{B_h : \text{temp_higher_}20^\circ}{I_h : \text{conditioning_on}} \mathcal{B2I} \quad \supset E
\end{array}$$

FIGURE 3. Deduction of $I_h : \text{conditioning_on}$, from $B_h : B_s \text{temp_higher_}20^\circ \supset \text{temp_higher_}20^\circ$, with axiom $B_h B_s : \neg \text{conditioning_on} \supset \text{temp_higher_}20^\circ$

$$\frac{B_h : \text{temp_higher_}20^\circ}{I_h : \text{conditioning_on}} \mathcal{B2I}. \quad (3.1)$$

If the hearer believes that the temperature is higher than 20 degrees, then it has the intention of switching conditioning on.

$$\text{AX} = \{B_h : (3.2), B_h : (3.3), B_h : (3.4), B_h B_s : (3.5)\}$$

where:

$$B_s \text{temp_higher_}20^\circ \supset \text{temp_higher_}20^\circ. \quad (3.2)$$

If s believes that the temperature is higher than 20 degrees, then the temperature is higher than 20 degrees.

$$\text{conditioning_on} \supset \neg \text{temp_higher_}20^\circ. \quad (3.3)$$

If the conditioning is on, then the temperature is lower than 20 degrees.

$$B_s (\text{temp_higher_}20^\circ \wedge \text{conditioning_on}) \supset I_s \text{stop_working}. \quad (3.4)$$

If s believes that the temperature is higher than 20 degrees and that the conditioning is on, then s intends to stop working.

$$\neg \text{conditioning_on} \supset \text{temp_higher_}20^\circ. \quad (3.5)$$

If the conditioning is off, then the temperature is higher than 20 degrees.

With the set of axioms AX, h infers that, if s believes that the conditioning is off, then h adopts the intention of switching it on. In symbols:

$$\text{AX}, B_h : B_s \neg \text{conditioning_on} \vdash_{\text{MC}} I_h : \text{conditioning_on}.$$

The corresponding deduction in MC is shown in Figure 3.

4 Plan-based model of speech acts

The plan-based vision of *speech acts* [5], which treats them as actions and represents them as STRIPS-like operators, offers us an intuitive way to correlate the speaker's mental state to its utterances. The 'trick' is in the modeling of the speech acts' preconditions. To illustrate this idea we use a simplified and revised version of the INFORM and REQUEST operators

which are, respectively, the prototypical members of the assertive and directive speech act classes [32].

In a plan-based theory of speech acts, $\text{INFORM}(s, h, \phi)$ is generally defined to be an action whose main effect on hearer's mental state is that the hearer believes that the speaker believes the propositional content ϕ , and its prerequisite is that the speaker believes ϕ (sincerity).

<i>Speech act</i>	<i>Preconditions</i>	<i>Main effects</i>
$\text{INFORM}(s, h, \phi)$	$\phi \in B_s^*$	$\phi \in B_h B_s^*$

The structure of this simple operator is closely related to the one by Cohen and Perrault [5]. We envisage, however, a larger range of effects described in the following:

1. The effects of the INFORM operator given by Cohen and Perrault [5] are the main effects here. The complete effects of a speech act on the hearer's mental state are beyond the speaker's control. We think that part of these *perlocutionary effects* are the result of some kind of abductive reasoning performed by the hearer from the received communication and from its actual mental state (which is in general different from the speaker's image of the hearer's mental state).
2. The precondition of the INFORM operator does not include the speaker's goal to perform such a speech act. As we see later, the speaker's actual intentions which leads to the execution of the speech act, are ascribed by the hearer to the speaker, again by some kind of abductive reasoning.

Note that the preconditions of the speech act must be verified in the speaker's mental state, while the effects must be verified both in the hearer's and the speaker's mental state.

Without loss of generality, in this paper we restrict the $\text{REQUEST}(s, h, \phi)$ operator to yes/no questions. This speech act is necessary when the speaker's beliefs are not sufficient to infer something that it is interested in. Hence, its preconditions are $\phi \notin B_s^*$ and $\neg\phi \notin B_s^*$. As in the case of INFORM, the effects are that the hearer believes the operator's preconditions, and that the speaker believes that the hearer believes them. $\text{REQUEST}(s, h, \phi)$ might also express the speaker's intention to believe ϕ rather than believe $\neg\phi$ (otherwise the speech act should have been $\text{REQUEST}(s, h, \neg\phi)$). This yields two further effects, namely that: ' h believes that s has the intention of believing ϕ ', and that ' s believes that h believes that s has the intention of believing ϕ '. Then we could define the REQUEST operator as:

<i>Speech act</i>	<i>Preconditions</i>	<i>Main effects</i>
$\text{REQUEST}(s, h, \phi)$	$\phi \notin B_s^*$ $\neg\phi \notin B_s^*$	$\phi \notin B_h B_s^*$ $\neg\phi \notin B_h B_s^*$ $\phi \notin B_s B_h B_s^*$ $\neg\phi \notin B_s B_h B_s^*$ $\phi \in B_h I_s B_s^*$ $\phi \in B_s B_h I_s B_s^*$

The basic assumption in this paper is that there is a causal relationship between an agent's mental state and its possible uttering a sentence. We may say that s plans an $\text{INFORM}(s, h, \phi)$ because of the facts that:

- II.** s has the intention of bringing h in a mental state where a formula ψ (which might differ from ϕ itself) is either believed or intended by h ; i.e. either $\psi \in I_s B_h^*$ or $\psi \in I_s I_h^*$.

- I2.** s does not believe that h is already in that mental state: $\psi \notin B_s B_h^*$ (resp. $\psi \notin B_s I_h^*$).
- I3.** s believes that if it performs the INFORM act, then h will be in a mental state in which it believes (resp. intends) ψ ; i.e. $\psi \in B_s B_h'^*$ (resp. $\psi \in B_s I_h'^*$).
- I4.** s can perform the INFORM; that is, the INFORM's precondition holds before and after performing the speech act; i.e. $\phi \in B_s^*$ and $\phi \in B_s'^*$.

We may say as well as that s plans a REQUEST(s, h, ϕ) because:

- R1.** s has the intention of being in a certain mental state in which it either believes or intends a certain formula ψ ; i.e. either $\psi \in I_s B_s^*$ or $\psi \in I_s I_s^*$.
- R2.** s does not believe that it is already in that mental state: $\psi \notin B_s B_s^*$ (resp. $\psi \notin B_s I_s^*$).
- R3.** s can perform the REQUEST act; that is, the REQUEST's preconditions hold before the performing of a REQUEST speech act; i.e. $\phi \notin B_s^*$ and $\neg\phi \notin B_s^*$.

5 Contextual abduction and revision

In this section, we first propose a formal framework that generalizes the ideas of abductive expansion to multi-context systems; and then, we introduce three basic operations for updating the hearer's mental state.

5.1 Multi-context abduction

Let briefly recall the main concepts of causal theory, abduction, and abductive explanation introduced by Konolige in [19]. Roughly speaking, abduction is an abstract hypothetical inferential schema that, given a causal theory of the domain, and an observation on a set of observable effects, looks for an explanation for them. An explanation for an observation is a minimal set of hypothesis, chosen among a set of possible causes, which if "added" to the causal theory, justify the observation. Syntactic propositional accounts of abduction formalize *causes* and *effects* as literals of a finite propositional language L , and the causal theory of the domain (*domain theory*) as a propositional theory of L .

A *simple causal theory* on a finite propositional language L , is a tuple $T = \langle C, E, \Sigma \rangle$, where C and E are sets of literals of L , and Σ is a theory on L . An *abductive explanation* (ABE) of an observation $O \subseteq E$, under a domain theory Σ , is a finite set $A \subseteq C$ such that:

- $\Sigma \cup A \not\vdash \perp$ (A is consistent with Σ),
- $\Sigma \cup A \vdash O$,
- A is subset-minimal over sets satisfying the first two conditions.

DEFINITION 5.1

A *simple multi-context causal theory* MT for a multi-context system MC is a family $\{T_\alpha = \langle C_\alpha, E_\alpha, \Sigma_\alpha \rangle\}$, where for each context α , T_α is a simple causal theory and $E_\alpha \subseteq C_\alpha$.

Introducing the hypothesis $E_\alpha \subseteq C_\alpha$, we accept the fact that each effect can be regarded as the explanation of itself.

DEFINITION 5.2

Let MT be a simple multi-context causal theory. The *causes* \mathbf{C} , the *effects* \mathbf{E} , and the *domain theory* Σ of MT are the sets of the formulae contained, respectively, in all C_α , E_α , and Σ_α of MT. In symbols, $\mathbf{C} = \{\alpha : \sigma \mid \sigma \in C_\alpha\}$, $\mathbf{E} = \{\alpha : \sigma \mid \sigma \in E_\alpha\}$, and $\Sigma = \{\alpha : \sigma \mid \sigma \in \Sigma_\alpha\}$.

DEFINITION 5.3

Let MT be a simple multi-context causal theory for a multi-context system MC. An *Abductive Explanation* (ABE) for an observation $O \subseteq E_\alpha$ in a context α under the domain theory Σ of MT, is a set $\mathbf{A} \subseteq \mathbf{C}$, such that:

1. $\forall \beta, \Sigma \cup \mathbf{A} \not\vdash_{\text{MC}} \beta : \perp$; \mathbf{A} is consistent with Σ in any context.
2. $\Sigma \cup \mathbf{A} \vdash_{\text{MC}} \alpha : O$; the observation O can be derived in α from the set of axioms \mathbf{A} and the domain theory Σ .
3. \mathbf{A} is the minimal set satisfying conditions 1 and 2. This means that for any set $\mathbf{B} \subseteq \mathbf{C}$ satisfying condition 1 and 2, $\mathbf{B} \not\subseteq \mathbf{A}$.

From the decidability of MC (see [15]), and the fact that the set of the causes is finite, we can conclude that the problem of finding an ABE of an observation O under the domain theory Σ is decidable.

EXAMPLE 5.4

Let consider the MC of the working example, and let suppose that the MT for it is composed by the domain theory $\Sigma = \{(3.2)\text{--}(3.5)\}$ and the sets of causes and effects defined as follows:

- $E_{B_h} = E_{B_h B_s} = E_{B_h B_s B_h} = \{\text{conditioning_on}, \text{temp_higher_}20^\circ\}$
Beliefs and nested beliefs about temperature and about the conditioning are considered observable effects.
- $E_{I_h} = E_{B_h I_s} = \{\text{conditioning_on}, \text{stop_working}\}$
Intentions and beliefs about the speaker's intentions regarding the conditioning and working are considered observable effects.
- For any other context α , $E_\alpha = \emptyset$
We are not interested in effects in contexts different from the ones mentioned above.
- $C_{B_h} = C_{B_h B_s} = C_{B_h B_s B_h} = \{\text{conditioning_on}, \text{temp_higher_}20^\circ\}$
Beliefs and nested beliefs about the temperature and about the conditioning are considered possible causes of the observable effects.
- $C_{I_h} = C_{B_h I_s} = \{\text{conditioning_on}, \text{stop_working}\}$
Beliefs and intentions regarding the speaker's intentions regarding the conditioning and working are considered possible causes of the observable effects.
- For any other context α , $C_\alpha = \emptyset$
We are not interested in causes in contexts different from the ones mentioned above.

An ABE for the observation $I_h : \text{conditioning_on}$ is $\{B_h B_s : \neg \text{conditioning_on}\}$. Indeed, we have that:

1. $\forall \beta$ of MC, $\Sigma, B_h B_s : \neg \text{conditioning_on} \not\vdash_{\text{MC}} \beta : \perp$;
2. $\Sigma, B_h B_s : \neg \text{conditioning_on} \vdash_{\text{MC}} I_h : \text{conditioning_on}$ (see deduction in Figure 3);
3. $\{B_h B_s : \neg \text{conditioning_on}\}$ is minimal.

An ABE of the observation $B_h I_s : \text{stop_working}$ is

$$\{B_h B_s : \text{temp_higher_}20^\circ, B_h B_s : \text{conditioning_on}\}. \quad (5.1)$$

Notice that $B_h B_s : \text{temp_higher_}20^\circ$ can be derived from $B_h B_s : \neg \text{conditioning_on}$ (by axiom (3.5)). This implies that the formula $B_h B_s : \text{temp_higher_}20^\circ$ could be replaced

by $B_h B_s : \neg \text{conditioning_on}$ in (5.1), to obtain a second ABE for the observation $B_h I_s : \text{stop_working}$. On the other hand, we cannot accept

$$\{B_h B_s : \neg \text{conditioning_on}, B_h B_s : \text{conditioning_on}\}$$

as an ABE since it violates the consistency condition.

5.2 Multi-context revision

When the hearer receives a message, it can do a number of observations on the conditions that have induced the speaker to send such a message. For instance, when h receives an INFORM from s , h can observe that the conditions **I1–I4** on the speaker's mental state, described in Section 4, must hold. Such observations, however, are not unconditionally accepted by the hearer, rather it looks for a set of explanations for them, and only if it finds some satisfactory explanations does it update its mental state accordingly.

In order to explain observations deriving from communication, the hearer must be provided with a simple multi-context causal theory MT for MC. The domain theory of MT must be always part of h 's mental state, and as the hearer's mental state is completely determined by the set of axioms AX (as described in Section 2), Σ must be a subset of AX (in symbols, $\Sigma \subseteq \text{AX}$). However, the domain theory must never be revised by the hearer, whereas AX can contain other revisable beliefs and intentions, which can change along the dialogue. As argued above, these changes are the result of accepting the explanations of some observation. As a consequence, the portion of AX, which is not in Σ , must be a subset of the causes; in symbols, $\text{AX} = \Sigma \cup \mathbf{X}$, where $\mathbf{X} \subseteq \mathbf{C}$ is the only part that can be modified by receipt of the speech act (i.e., $\text{AX}' = \Sigma \cup \mathbf{X}'$). We call \mathbf{X} the set of the *current explanations*. \mathbf{X} can be changed by applying three basic operations: *abductive expansion*, *abductive contraction*, and *abductive revision*.

Abductive expansion Abductive expansion, denoted by $+$, is applied when \mathbf{X} must be extended to derive an observation ϕ in a context α :

$$\mathbf{X} + \alpha : \phi \triangleq \mathbf{X} \cup \mathbf{A}$$

where \mathbf{A} is an ABE of $\alpha : \phi$ under the domain theory $\Sigma \cup \mathbf{X}$. Expansion does not specify how \mathbf{A} is chosen among the set of the minimal explanations. This choice might be based on a partial order on ABEs. In this paper we do not consider the effect and the specific definition of such an order. This order strictly depends on the application domain and on the meaning and the degree of plausibility of the explanations. Similarly, we do not consider methods to represent and compute such a partial order.

Abductive contraction The operation of contraction, denoted by $-$, is applied if \mathbf{X} is in contrast with a new observation ϕ , and then some formulae must be removed from \mathbf{X} . For any formula ϕ in a context α :

$$\mathbf{X} - \alpha : \phi \triangleq \mathbf{X} \setminus \mathbf{Y}$$

where \mathbf{Y} is a minimal hitting set² of $\{\mathbf{A}_1, \dots, \mathbf{A}_n\}$, and $\{\mathbf{A}_1, \dots, \mathbf{A}_n\}$ is the set of all the ABEs of ϕ in the context α under the domain theory Σ , which are contained in \mathbf{X} . Again, we

²Given a collection of sets $S = \{\mathbf{A}_1, \dots, \mathbf{A}_n\}$, a set \mathbf{H} is a hitting set of S if for each \mathbf{A}_i , $\mathbf{H} \cap \mathbf{A}_i \neq \emptyset$. A hitting set is minimal if for any other hitting set \mathbf{H}' , $\mathbf{H}' \not\subseteq \mathbf{H}$.

have not specified how the hitting set is chosen. As before, this will be related to the ordering on the ABEs.

Abductive revision The operation of revision, denoted by $*$, is applied when \mathbf{X} explains something which is inconsistent with an observation ϕ in a context α :

$$\mathbf{X} * \alpha : \phi \triangleq (\mathbf{X} - \alpha : \neg\phi) + \alpha : \phi.$$

This operative definition comes from the Levi identity [9, 7].

6 Updating mental states from communication

In this section we formally characterize the methods for updating the hearer's image of the speaker's mental state as a consequence of receipt of the speech acts INFORM and REQUEST.

6.1 Updating from INFORM

Checking preconditions (condition I4) Suppose that s sends to h a message of the form $\text{INFORM}(s, h, \phi)$. Being aware of the fact that condition **I4** holds, h may update its image of the speaker's mental state by imposing that the precondition of $\text{INFORM}(s, h, \phi)$ holds on its image of the speaker's beliefs³:

$$\phi \in B_h B_s^*.$$

The new mental state is obtained by updating the set of the current explanations \mathbf{X} as follows:

$$\mathbf{X}' = \mathbf{X} * B_h B_s : \phi.$$

We can distinguish two cases: either $\Sigma \cup \mathbf{X}$ is consistent with $B_h B_s : \phi$, (i.e. $\neg\phi \notin B_h B_s^*$) or it is inconsistent with $B_h B_s : \phi$ (i.e., $\neg\phi \in B_h B_s^*$). In the first case, h simply expands \mathbf{X} with an explanation of $B_h B_s : \phi$. In the second case, h computes \mathbf{X}' first by contracting \mathbf{X} , in order to have a new set $\mathbf{Y} = \mathbf{X} - B_h B_s : \neg\phi$ such that $\Sigma \cup \mathbf{Y} \not\vdash_{\text{MC}} B_h B_s : \neg\phi$, and then expanding \mathbf{Y} into a set $\mathbf{X}' = \mathbf{Y} + B_h B_s : \phi$ by adding an explanation of $B_h B_s : \phi$. In the resulting mental state, we have that $\phi \in B_h B_s^*$.

EXAMPLE 6.1

Suppose that s sends to h the message $\text{INFORM}(s, h, \text{temp_higher_}20^\circ)$ when the set of h 's current explanations $\mathbf{X} = \emptyset$. It is easy to see that the new set of explanations \mathbf{X}' is equal to $\mathbf{X} * B_h B_s : \text{temp_higher_}20^\circ = \{B_h B_s : \neg\text{conditioning_on}\}$. Indeed, since $\neg\text{temp_higher_}20^\circ \notin B_h B_s^*$ the hearer computes the following minimal ABE:

$$\mathbf{A} = \{B_h B_s : \neg\text{conditioning_on}\}$$

and expands \mathbf{X} accordingly, resulting $\mathbf{X}' = \{B_h B_s : \neg\text{conditioning_on}\}$. Notice that, in this new mental state, h has the intention of switching the conditioning on; in fact, we have that $\Sigma \cup \mathbf{X}' \vdash_{\text{MC}} I_h : \text{conditioning_on}$.

³Notice that this coincides also with the main effects of $\text{INFORM}(s, h, \phi)$ on the mental state of h .

Intention recognition (condition I1) By intention recognition we mean the hearer's ability to recognize the intention that induced the speaker to perform the speech act. Condition **I1** states that a motivation for s to perform an $\text{INFORM}(s, h, \phi)$ is its intention of changing h 's mental state so that h believes or intends some new formula. To discover this intention, h checks the differences between its mental state before and after s executes $\text{INFORM}(s, h, \phi)$ (\mathbf{X} and \mathbf{X}' respectively), and then it revises \mathbf{X}' to include the fact that s has the intention of causing those differences. For instance, suppose that for the context B_h ⁴ there is a formula ψ such that:

$$\psi \notin B_h^* \text{ and } \psi \in B_h'^* \quad (6.1)$$

Intuitively, this means that one of the effects that the speaker has obtained by its utterance, is that the hearer believes ψ . Supposing that this is the intention of the speaker, the hearer can revise \mathbf{X}' by the observation $B_h I_s B_h : \psi$; namely, h can try to find an explanation of the fact that the speaker has the intention of making h believe ψ . The new set of the current explanations \mathbf{X}'' is obtained as follows:

$$\mathbf{X}'' = \mathbf{X}' * B_h I_s B_h : \psi.$$

Similar revision can be done on any other context of the hearer's mental state, considering the effects obtained by the speaker's utterance in that context.

EXAMPLE 6.2

Let restrict the intention recognition problem to the possible speaker's intentions about the hearer's beliefs and intentions. Suppose that s sends to h the message $\text{INFORM}(s, h, \text{temp_higher_}20^\circ)$ and that \mathbf{X}' is computed as in Example 6.1. In this case, h finds that $\text{temp_higher_}20^\circ$ is not in B_h^* but it is in $B_h'^*$. Supposing that this is the intention of the speaker, h can revise \mathbf{X}' as follows:

$$\mathbf{X}'' = \mathbf{X}' * B_h I_s B_h : \text{temp_higher_}20^\circ.$$

Moreover, h finds that $I_h : \text{conditioning_on}$ is a formula that does not belong to \mathbf{X} , but belongs to \mathbf{X}' . As before, supposing that the speaker has the intention of making h intend to bring about conditioning_on , h can revise \mathbf{X}' as follows:

$$\mathbf{X}'' = \mathbf{X}' * B_h I_s I_h : \text{conditioning_on}.$$

A final update (condition I3) If the hearer has recognized the speaker's intentions, then it can update its image of the speaker's mental state in order to have the speaker believe that its intentions have been satisfied (condition **I3**). In particular, in the case (6.1), h can revise \mathbf{X}'' in order to verify that s believes that h believes ψ ; in symbols:

$$\mathbf{X}''' = \mathbf{X}'' * B_h B_s B_h : \psi.$$

6.2 Updating from REQUEST

Checking preconditions (condition R3) Let us suppose that s sends to h the message $\text{REQUEST}(s, h, \phi)$. Being aware of the fact that condition **R3** holds, h may update its image of the speaker's mental state by imposing that the preconditions of $\text{REQUEST}(s, h, \phi)$

⁴ h could be interested only in the effects yielded in some particular context.

hold on its image of the speaker's beliefs⁵:

$$\phi \notin B_h B_s^* \text{ and } \neg\phi \notin B_h B_s^*.$$

The new mental state is obtained by contracting the set of the current explanations \mathbf{X} as follows:

$$\mathbf{X}' = \mathbf{X} - B_h B_s : \phi \text{ or } \mathbf{X}' = \mathbf{X} - B_h B_s : \neg\phi.$$

EXAMPLE 6.3

Suppose that s performs $\text{REQUEST}(s, h, \text{conditioning_on})$ when the set of the current explanations is \mathbf{X}'' as computed in Example 6.2. Some precondition of REQUEST is not verified, as $B_h B_s : \neg\text{conditioning_on}$ is in \mathbf{X}'' . Then, h can update its image of s ' mental state by contracting \mathbf{X}'' as follows:

$$\mathbf{X}''' = \mathbf{X}'' - B_h B_s : \neg\text{conditioning_on}.$$

Intention recognition (conditions R1 and R2) Condition **R1** states that a motivation for s to perform a $\text{REQUEST}(s, h, \phi)$ is its intention of being in a certain mental state in which it either believes or intends some new formula. To discover this intention, h checks the differences between its image of s ' mental state before and after s believes in the REQUEST 's content (\mathbf{X} and \mathbf{X}' respectively), and then it revises \mathbf{X}' to include the fact that s has the intention of causing those differences. For instance, suppose that for the context $B_h I_s$ there is a formula ψ such that:

$$\psi \notin B_h I_s^* \text{ and } \psi \in B_h I_s'^*.$$

This means that one of the effects that the speaker has obtained believing the content of its request, is that now it intends ψ . Supposing that this is the intention of the speaker, the hearer can revise \mathbf{X}' by the observation $B_h I_s : \psi$; namely, h can try to find an explanation of the fact that the speaker has the intention of bringing about ψ . The new set of current explanations \mathbf{X}'' is obtained as follows:

$$\mathbf{X}'' = \mathbf{X}' * B_h I_s : \psi.$$

In general, if the effects are observed in a context α , h revises \mathbf{X}' as follow:

$$\mathbf{X}'' = \mathbf{X}' * B_h I_s \alpha : \psi.$$

EXAMPLE 6.4

Let suppose that s performs $\text{REQUEST}(s, h, \text{conditioning_on})$ and that the set of h 's current explanations is \mathbf{X}''' as computed in Example 6.3. In this case, h finds that stop_working is not in $B_h I_s^*$ but it is in $B_h I_s'^*$. Supposing that this is the intention of the speaker, h revises \mathbf{X}''' as follows:

$$\mathbf{X}'''' = \mathbf{X}''' * B_h I_s : \text{stop_working}.$$

This means that h believes that s intends stop_working .

⁵Notice that this coincides also with the main effects of $\text{REQUEST}(s, h, \phi)$ on the mental state of h .

7 Related work

The work presented in this paper relates with four main research areas: agents' internal structure, agent's communication language, abduction, and belief revision.

Concerning agents' internal structure, we based our approach on some previous well established results. Agents' mental states are formalized by multi-context systems, analogously to what has been done in [15, 3], and planning of communication has been formalized by extending the plan-based theory of speech acts described in [5]. In [26], Parsons *et al.* propose a formal model of argumentation-based reasoning and negotiation for autonomous agents. Although, they use multi-context systems to model BDI agents, in their framework the agents have no explicit representations of the other agents' mental state. The main reason of this is that the agents exchange part of their mental state during the negotiation. Differently, in our work we suppose that the agents communicate one another without arguing the content of their communication.

In the literature, the major agent communication languages are: Knowledge Query and Manipulation Language (KQML) [8] and FIPA ACL [1] (see [20] for a detailed evaluation and comparison). Both are based on the classical speech act theory and differ primarily in the details of their semantic framework. KQML and FIPA ACL do not specify how their semantics must be used by the agents involved in a dialogue. Differently, in our approach we have defined a concrete and computationally feasible way for an agent to treat a specific set of communicative acts. In other words, given a set of communicative acts with semantics expressed in terms of preconditions and main effects, we have defined a set of updating policies that an agent can follow to update its mental state whenever it receives a message.

In spite of the two well-grounded research areas on 'belief revision' and 'abduction', few works attempted to combine them in a unified treatment. In [21], Lobo and Uzcátegui define an abductive version of a large class of theory change operators. For any operator $*$, they define its abductive version $*_a$. In particular, they define abductive expansion, abductive revision, abductive contraction, and abductive update on the basis of their respective classical versions defined in the literature. Although this approach seems to share many common aspects with the one presented in this paper, it differs in the motivations and in the final goal. Their objective is to define general abductive change operators, whereas our work has the main goal to define a specific set of theory change operators suitable for observations generated by a communicative act. One of the main consequences of this difference concerns the definition of explanation set. We update (revise) our theory with respect to a single explanation chosen among the set of the minimal explanations, whereas they adopt the cautious explanation, that is the disjunction of all the minimal explanations. Our choice is motivated by the fact that we believe that in specific application domains, it is possible to define a partial order on the ABEs, and then choose the best explanation for a given observation.

Aravidian and Dung, in [2], state a number of rationality postulates for the contraction of a knowledge base with respect to a sentence, and they define an abduction based algorithm for its computation. Their algorithm (based on hitting sets) is very similar to the definition of abductive contraction given in this paper. As a matter of fact, our definition of contraction fulfils their basic rationality postulates. A second important analogy is the fact they suppose that the knowledge base is composed of two subsets: an 'immutable theory' and an 'updatable theory', which are the analogous of Σ and X defined in this paper. The main difference is that we extend this idea to the case of abductive expansion and revision, and furthermore, that we specialize the operators to a logic for beliefs and intentions.

Analogously, Pagnucco *et al.* [24, 25] introduce some rationality postulates for abductive

expansion. In [23], they argue that the notion of abduction corresponds to an attempt to determine an initial belief state from a contracted belief state and an epistemic input, under certain conditions. In their work, the revision process is limited to the agent's beliefs. Introducing mental states, we extend the revision process to the agent's intentions. This is very important for multi-agent systems because it allows an agent to revise its intentions whenever new beliefs are acquired. Moreover, the use of image of mental states allows an agent to maintain its beliefs about the mental state of the other agents always updated.

Hindriks *et al.* in [18] provide an operational semantics, based on transition systems, for two pairs of communicative operators: ask and tell, and request and offer. As in our approach, each agent has a mental state composed of two sets: beliefs and goals. Additionally, each mental state contains a set of rules describing its evolution and it refers to a particular moment during the agent evolution. The semantics of communicative operators is given in terms of a transition of an agent from a belief state to another. A first difference between their and our approach is that they do not allow agents to have images of the other agents' mental states. A more radical difference regards the fact that in their semantics, communicative acts have neither preconditions nor effects on the beliefs of an agent. In other words, the semantics does not contain an explicit relationship between the agent's beliefs and its communication. The main consequence of this is that the receipt of a message does not necessarily yield the revision of the agent's mental state. Belief revision is considered as any other action. An agent can decide to revise its beliefs independently from the communication with other agents. Moreover, in [18] the authors do not provide any criteria for belief revision, which is assumed to be a pre-compiled function. Our proposal is therefore complementary to this approach as we provide some well founded and executable methods for revising beliefs after communication. Finally, Hindriks *et al.* deal with synchronous communication, whereas we consider asynchronous communication. Synchronous communication means that a tell from agent *A* to agent *B* is necessarily preceded by an ask from agent *B* to agent *A*, and vice versa any ask from *B* to *A* is followed by a tell in the opposite direction. This is true also for the other pair of communicative acts, namely request and offer. Deduction and abduction are not used to enlarge and/or modify agent's beliefs. Rather, deduction is a reasoning pattern which is used by an agent to generate an answer (tell) of an information request (ask); and similarly, abduction is used to generate an offer (offer) that fulfils a request (request). Differently, in our approach deduction and abduction are performed not only to generate the answer to a specific request, but to process the receipt and the sending of any message.

8 Conclusion

Communication can be effectively improved when agents are able to reason on other agents' mental state. In this paper we have made the fundamental assumption that there is a causal relationship between a speaker's mental state and its uttering a sentence. Following this idea, we have developed a set of criteria that allow the hearer to recognize and update its representation of the speaker's mental state on the basis of the speaker's utterance.

Inspiration for the causal relationship has been taken from the classical 'Speech Acts Theory'; namely, we have adopted the plan-based vision of speech acts in representing them as STRIPS-like operators. The two major agent communication languages based on the speech act theory (namely KQML and FIPA ACL) do not specify how their semantics must be used by the agents involved in the dialogue. In our approach we have overcome this difficulty defining a concrete method that a hearer can adopt to use both the preconditions and the

effects of the received speech act for updating its image of the speaker's mental state.

Our work is based upon the use of multi-context systems for which we have extended the notion of casual theories, abduction and revision. Combining abduction and belief revision in a unified treatment, we have defined three basic operations for updating. These operations are defined for multi-context systems and their properties are different from those of the classical AGM operators. In this paper we have not discussed these differences. Moreover, although decidability is guaranteed by the decidability of multi-context systems, we have not discussed the computational complexity of our approach. The study of the properties and the computational complexity of the operations for updating will be the object of future work.

References

- [1] Foundation for intelligent physical agents. ACL Specification version 2.0, <http://www.fipa.org/spec/fipa97.html>, 1997.
- [2] C. Aravindan and P.M. Dung. Knowledge revision, abduction and database updates. *Journal of Applied Non-Classical Logics*, **5**, 51–76, 1995.
- [3] 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**, 401–423, 1998.
- [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)*, pp. 57–64, 1996.
- [5] P.R. Cohen and C.R. Perrault. Elements of a plan-based theory of speech acts. *Cognitive Science*, **3**, 177–212, 1979.
- [6] A. F. Dragoni and P. Puliti. Mental states recognition from speech acts through abduction. In *Proceedings of the 11th European Conference on Artificial Intelligence*. John Wiley & Sons, 1994.
- [7] A.F. Dragoni and P. Giorgini. Revising beliefs received from multiple source. In *Frontiers of Belief Revision*, M. A. Williams and H. Rott, eds., Applied Logic. Kluwer, 2000.
- [8] T. Finin, R. Fritzson, D. McKay, and R. McEntire. KQML as an agent communication language. In *Proceedings of the Third International Conference on Information and Knowledge Management (CIKM'94)*, ACM Press, 1994.
- [9] P. Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. MIT Press., Cambridge, MA, 1988.
- [10] M. Georgeff. Communication and interaction in multiagent planning. In *Proceedings of the 3th National Conference on Artificial Intelligence*, pp. 125–129, 1983.
- [11] C. Ghidini. Modelling (un)bounded beliefs. In *Modelling and Using Context – Proceedings of the 2nd International and Interdisciplinary Conference, Context'99*, P. Bouquet, L. Serafini, P. Brezillon, M. Benerecetti, and F. Castellani, eds. Volume 1688 of *LNAI*, pp. 145–158. Springer Verlag, Heidelberg, 1999.
- [12] E. Giunchiglia and F. Giunchiglia. Ideal and real belief about belief. *Journal of Logic and Computation*, **11**, 2001. To appear.
- [13] F. Giunchiglia. Contextual reasoning. *Epistemologia, Special Issue on I Linguaggi e le Macchine*, XVI:345–364, 1993.
- [14] 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)*, pp. 282–289. Morgan Kaufmann, 1998.
- [15] F. Giunchiglia and L. Serafini. Multilanguage hierarchical logics (or: how we can do without modal logics). *Artificial Intelligence*, **65**, 29–70, 1994.
- [16] P.J. Gmytrasiewicz and E.H. Durfee. A rigorous, operational formalization of recursive modeling. In *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS)*, pp. 125–132, 1995.
- [17] P.J. Gmytrasiewicz, E.H. Durfee, and D.K. Wehe. The utility of communication in coordinating intelligent agents. In *Proceedings of the Ninth National Conference on Artificial Intelligence*, 1991.
- [18] K.V. Hindriks, F.S. de Boer, W. van der Hoek, and J.J.Ch. Meyer. Semantics of communicating agents based on deduction and abduction. In *Proceedings of IJCAI'99 Workshop on ACL*, 1999.
- [19] K. Konolige. Abduction versus closure in causal theories. *Artificial Intelligence*, **53**, 255–272, 1992.

- [20] Y. Labrou, T. Finin, and Y. Peng. Agent communication languages: The current landscape. *Intelligent Systems*, **14**, 1999.
- [21] J. Lobo and C. Uzcátegui. Abductive change operators. *Fundamenta Informaticae*, **27**, 319–418, 1996.
- [22] J.S. Rosenschein M.R. Gensereth, M.L. Ginsberg. Cooperation without communication. In *AAAI 86*, pp. 51–57, 1986.
- [23] M. Pagnucco and N.Y. Foo. The relationship between abduction and changes in belief states. In *Proceedings of the ICLP93 Postconference Workshop on Abductive Reasoning*, pp. 75–83, Budapest, Hungary, 1993.
- [24] M. Pagnucco, A.C. Nayak, and N.Y. Foo. Abductive expansion: The application of abductive inference to the process of belief change. In *Proceedings of the Seventh Australian Joint Conference on Artificial Intelligence*, pp. 70–77, Armidale, Australia, 1994.
- [25] M. Pagnucco, A.C. Nayak, and N.Y. Foo. Abductive reasoning, belief expansion and nonmonotonic consequence. In *Proceedings of the ICLP'95 Joint Workshop on Deductive Databases and Logic Programming and Abduction in Deductive Databases and Knowledge-based Systems*, pp. 143–158, Shonan Village Center, Japan, 1995.
- [26] S. Parsons, C. Sierra, and N.R. Jennings. Agents that reason and negotiate by arguing. *Journal of Logic and Computation*, **8**, 529–577, 1998.
- [27] D. Prawitz. *Natural Deduction - A Proof Theoretical Study*. Almquist and Wiksell, Stockholm, 1965.
- [28] 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)*, pp. 312–319, S. Francisco, CA, 1995.
- [29] A.S. Rao, M. Georgeff, and E.A. Sonenberg. Social plans: A preliminary report. In *Decentralized AI - Proceedings of the Third European Workshop on Modeling Autonomous Agents in a Multi-Agent World (MAAMAW-91)*, E. Werner and Y. Demazeau, eds. pp. 57–76, Elsevier Science Publishers B.V., Amsterdam, 1992.
- [30] A.S. Rao and G. Murray. Multi-agent mental-state recognition and its application to air-combat modelling. In *Proceedings of the Distributed Artificial Intelligence Workshop*, Lake Quinalt, Washington, 1994.
- [31] J.S. Rosenschein and M.R. Genesereth. Communication and cooperation. Stanford Heuristic Programming Report, 1984.
- [32] J.R. Searle and D. Vanderveken. *Foundations of Illuccionary Logic*. Cambridge University Press, 1985.
- [33] 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*, Morgan Kaufmann, Los Altos, CA, 1988.

Received 8 August 1997