UNIVERSITY
OF TRENTO

DEPARTMENT OF INFORMATION AND COMMUNICATION TECHNOLOGY

38050 Povo – Trento (Italy), Via Sommarive 14
http://www.dit.unitn.it

WHAT IS LOCAL MODELS SEMANTICS?

Chiara Ghidini and Fausto Giunchiglia

November 2002

Technical Report # DIT-02-0090

Also: to appear as a chapter P. Bouquet, L. Serafini, "Perspectives on context", CSLI

What is Local Models Semantics?*

Chiara Ghidini[1] Fausto Giunchiglia[2]

[1] Department. of Computer Science, University of Liverpool
Liverpool L69 7ZF, United Kingdom

[2] Department of Information and Communication Technology
Università degli Studi di Trento. 38050 Povo, Trento, Italy.

chiara@csc.liv.ac.uk fausto@dit.unitn.it

Abstract

In recent papers a new semantics, called *Local Models Semantics*, was presented and used to provide a foundation to reasoning with contexts. Local Models Semantics captures and makes precise the two main principles underlying contextual reasoning: the, so-called, Principle of Locality and Principle of Compatibility. In this chapter we aim at explaining the main intuitions underlying Local Models Semantics, its fundamental logical properties, and its relation with contextual reasoning. The emphasis is on motivations and intuitions, rather than on technicalities.

1 Introduction

In recent papers a new semantics, called *Local Models Semantics*, was presented and used to provide a foundation to reasoning with contexts. An exhaustive presentation of the notion of context is out of the scope of this chapter.[1] The notion of context we consider here is based on two significative (informal) definitions independently proposed by Fausto Giunchiglia [14] and John McCarthy [20] in the late 80's, when context was introduced as an important means for formalising certain forms of reasoning.

According to [14], contexts are a tool for formalising the locality of reasoning:

> Our intuition is that reasoning is usually performed on a subset of the global knowledge base. The notion of context is used as a means of formalising this idea of localisation. Roughly speaking, we take a context to be the set of facts used locally to prove a given goal plus the inference routines used to reason about them (which in general are different for different sets of facts) [14].

In [20], contexts are introduced as a means for solving the problem of generality:

> When we take the logic approach to AI, lack of generality shows up in that the axioms we devise to express common sense knowledge are too restricted in their applicability for a general common sense database [...] Whenever we write an axiom, a critic can say that the axiom is true only in a certain context. With a little ingenuity the critic can usually devise a more general context in which the precise form of the axiom doesn't hold. [19]

Coherently with these two proposals, contexts have been used in various applications and in different domains. Contexts are used to deal with issues concerning the integration of heterogeneous knowledge and data bases. See for instance [7; 21; 12; 24]. The largest common-sense knowledge-base, CYC [18], contains an explicit notion of context [17]. Several references can be found in the literature about the use of contexts in the formalisation of reasoning about beliefs, meta reasoning, and propositional attitudes.

*Most of the material presented in this paper is based on the article [Ghidini and Giunchiglia, 2001] with title "Local Models Semantics, or Contextual Reasoning = Locality + Compatibility".

[1] The interested reader may refer to [1; 10; 14] for an accurate discussion on this topic.

See for instance [16; 13; 3; 8; 9; 15]. In [2] contexts are introduced in the formalisation of reasoning with viewpoints. [5] addresses the problem of formalising context-based common-sense reasoning. Finally, [4; 6; 23; 22; 11] introduce contexts to model different aspects of agents and multi-agent systems.

In spite of the variety of different approaches, formalisations, and domains of application, in [10] the authors claim that there are two main intuitions underlying the use of context, and state them as the following two principles:

Principle 1 (of Locality): reasoning uses only part of what is potentially available (e.g., what is known, the available inference procedures). The part being used while reasoning is what we call *context* (of reasoning);

Principle 2 (of Compatibility): there is *compatibility* among the reasoning performed in different contexts.

Local Models Semantics provides a formal framework where the two principles of Locality and Compatibility are captured and made precise. The goal of this chapter is to explain the main intuitions underlying Local Models Semantics, its fundamental logical properties, and its relations with contextual reasoning. The emphasis is on motivations and intuitions, rather than on technicalities. The reader interested in a more technical presentation and a detailed comparison with other logical frameworks may refer to [10]

The chapter is organised as follows. The core definitions are given in Sections 3 and Section 4. In Section 5 we comment on the properties of Local Models Semantics. In particular we investigate how the notion of context is formally defined within Local Models Semantics, and how Local Models Semantics captures the principles of Locality and Compatibility introduced above. In Section 6 we comment on how Local Models Semantics is able to deal with situations where we may or may not have a complete description of the world. To make the presentation clearer, in Section 2 we introduce a simple example of reasoning with viewpoints, called the *magic box* example, which will be used throughout the chapter. This example is a variation of the one originally proposed in [10].

2 The magic box example

Suppose there are two observers, Mr. Blue and Mr. Pink, each having a partial view of a box as shown in Figure 1. The box is composed of six sectors, each sector possibly containing a ball. There must be exactly

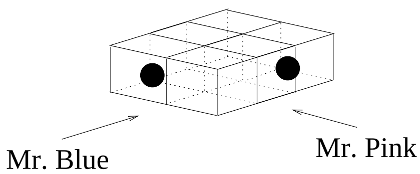

Figure 1: The magic box.

two balls in the box and there cannot be balls hidden from the view of an observer. The box is "magic" and observers cannot distinguish the depth inside it. Figure 2 shows the views of Mr. Blue and Mr. Pink corresponding to the scenario depicted in Figure 1.

In this example we focus on the two contexts describing the viewpoints of the two observers and the consequences that they are able to draw from it. The content of the two contexts corresponding to the scenario depicted in Figure 1 is graphically represented in Figure 2.

It is easy to see that the notions of locality and compatibility play a central role in this example. First locality. Both Mr. Blue and Mr. Pink have the notions of a ball being on the right or on the left. However we may have a ball which is on the right for Mr. Blue and not on the right for Mr. Pink. Furthermore Mr. Pink has the notion of "a ball being in the center of the box" which is meaningless for Mr. Blue. We also assume that the box is made of different coloured glass. Different observers, looking at the box from different sides,

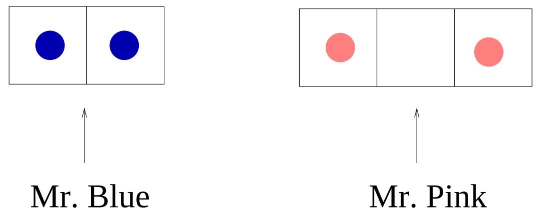

Figure 2: The contexts of Mr. Blue and Mr. Pink.

see the balls as if in different colours. In our example Mr. Blue sees (has the notion of) a ball being blue, while Mr. Pink sees (has the notion of) a ball being pink.

Focusing on compatibility, the contents of the contexts of Mr. Blue and Mr. Pink are obviously related. The relation is a consequence of the fact that Mr. Blue and Mr. Pink see the same box. Given the fact

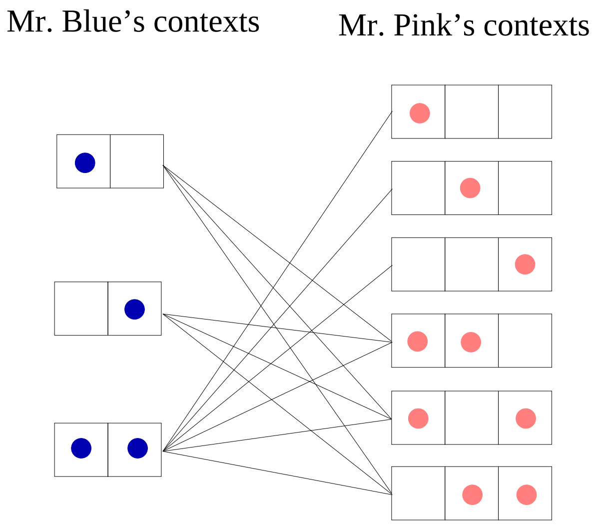

Figure 3: Compatible contexts of Mr. Blue and Mr. Pink.

that there must be exactly two balls in the box, it is easy to see that if Mr. Blue sees only one blue ball in the box, then Mr. Pink must see two pink balls in the box. Therefore we can describe this situation by listing all the possible compatible pairs (as they are represented in Figure 3), or we can describe it more synthetically using descriptions like: “if Mr. Blue sees a single blue ball then Mr. Pink sees two pink balls” and “if Mr. Pink sees a single pink ball then Mr. Pink sees two blue balls”.

3 Local models and model

We begin here the presentation of Local Models Semantics by defining the notions of local model and model.

3.1 The formal definitions

Let $\{L_i\}_{i\in I}$ be a family of languages defined over a set of indexes I (in the following we drop the index $i \in I$). Intuitively, each L_i is the formal language used to describe what is true in a context. For the purpose of our work we suppose that I is at most countable and that $\{L_i\}$ is a class of propositional languages. The first step towards the definition of a model for $\{L_i\}$ is to consider the class of models for each language L_i in

$\{L_i\}$. This will ensure that each language L_i is interpreted in its own, possibly different, structure. Formally, we denote with $\overline{M}_i$ the class of all the models of L_i. We call $m \in \overline{M}_i$ a *local model* (of L_i).

Then, we have to pair local models into a single structure. This is done by introducing the notions of compatibility sequence and compatibility relation. Formally, a *compatibility sequence* $\mathbf{c}$ (for $\{L_i\}$) is a sequence

$$\mathbf{c} = \langle \mathbf{c}_0, \mathbf{c}_1, \ldots, \mathbf{c}_i, \ldots \rangle$$

where, for each $i \in I$, $\mathbf{c}_i$ is a subset of $\overline{M}_i$. We call $\mathbf{c}_i$ the i-th element of $\mathbf{c}$. If $I = \{1, 2\}$ is composed of two indexes, a compatibility sequence $\mathbf{c}$ is of the form $\mathbf{c} = \langle \mathbf{c}_1, \mathbf{c}_2 \rangle$ and is called a *compatibility pair.*

A *compatibility relation* $\mathbf{C}$ (for $\{L_i\}$) is a set $\mathbf{C} = \{\mathbf{c}\}$ of compatibility sequences $\mathbf{c}$.[2]

We define a model as a compatibility relation which contains at least one sequence and does not contain the sequence of empty sets. Formally, a *model* (for $\{L_i\}$) is a compatibility relation $\mathbf{C}$ such that:

1. $\mathbf{C} \neq \emptyset$;
2. $\langle \emptyset, \emptyset, \ldots, \emptyset, \ldots \rangle \notin \mathbf{C}$.

In the following we write $\mathbf{C}$ to mean either a compatibility relation or a model, the context always makes clear what we mean.

In a nutshell, we can split the construction we perform into three steps. First, we start with some language, say L_1, L_2, and L_3 (see Figure 4). Then, we associate each L_i with a set $M_i \subseteq \overline{M}_i$ of local models. Usually $M_i \subset \overline{M}_i$ (see Figure 5). Finally, we pair local models inside compatibility sequences. The resulting compatibility relation is our model (see Figure 6).[3] Local models describe what is locally true. Compatibility sequences put together local models which are "mutually compatible", consistently with the situation we are describing. What we obtain are models composed of sets of "mutually compatible" sequences of local models.

Given a family of languages $\{L_i\}$, different classes of models may be defined, depending on the definition of compatibility relation. Different compatibility relations model different situations. A general class of models which will be used often in the chapter is based on the notion of chain. A compatibility sequence $\mathbf{c}$ is a *chain* if all the $\mathbf{c}_i$ contain exactly one local model (formally, if $|\mathbf{c}_i| = 1$ for each $i \in I$). A model $\mathbf{C}$ is a *chain model* if all the $\mathbf{c}$ in $\mathbf{C}$ are chains.

3.2 A model for the magic box

Let us apply the three step construction of the model depicted in Figures 4, 5, and 6 to the magic box example.

Languages We define the propositional languages L_B and L_P used by Mr. Blue and Mr. Pink, respectively, to describe their views. Let $P_B = \{r, l\}$ and $P_P = \{r, c, l\}$ be two sets of propositional constants. Intuitively, r, c, l stand for ball on the right, in the center and on the left, respectively. L_B is formally defined as the smallest set containing P_P, the symbol for falsity $\bot$, and closed under implication; L_P is formally defined as the smallest set containing P_2, the symbol for falsity $\bot$ and closed under implication. In this chapter we use the standard abbreviations from propositional logic, such as $\neg\phi$ for $\phi \supset \bot$, $\phi \vee \psi$ for $\neg\phi \supset \psi$, $\phi \wedge \psi$ for $\neg(\neg\phi \vee \neg\psi)$, $\top$ for $\bot \supset \bot$.

Local models We construct all the possible situations (local models) for L_B and L_P. L_B and L_P have the usual propositional semantics. Therefore the local models of L_B and L_P are univocally defined by sets of propositional formulae. In particular, the local models of L_B are univocally denoted by the following sets of formulae:

$$m_1 = \{l\} \qquad m_2 = \{r\} \qquad m_3 = \{l, r\}$$

[2] Formally, let $\prod_{i \in I} 2^{\overline{M}_i}$ be the Cartesian product of the collection $\{2^{\overline{M}_i} : i \in I\}$. The compatibility relation $\mathbf{C}$ is a relation of type $\mathbf{C} \subseteq \prod_{i \in I} 2^{\overline{M}_i}$

[3] Figures 4, 5, and 6 first appeared in [10].

Figure 4: Languages: L_1, L_2, and L_3.

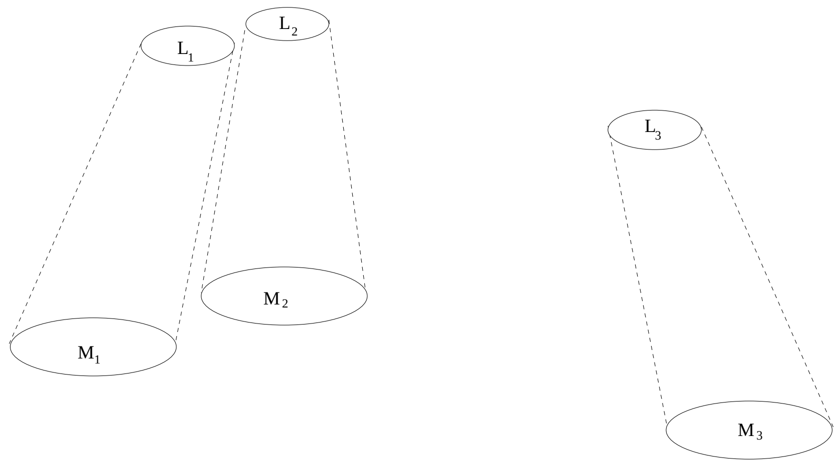

Figure 5: Local models for L_1, L_2, and L_3.

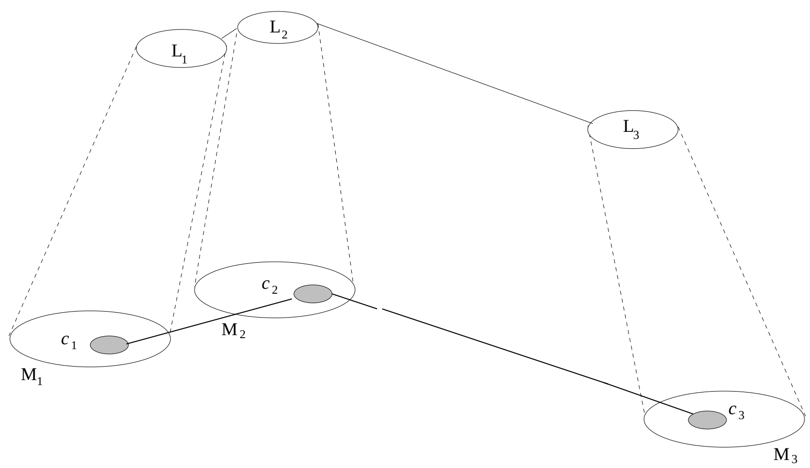

Figure 6: Model for $\{L_1, L_2, L_3\}$.

where we write $\{l\}$ to mean the local model describing the situation with a ball on the left, $\{r\}$ to mean the local model describing the situation with a ball on the right, and $\{l, r\}$ describing the situation with a ball on the left and a ball on the right.

Analogously, the local models of L_P are univocally denoted by the following sets of formulae:

$$\begin{array}{lll} m_1 = \{l\} & m_2 = \{c\} & m_3 = \{r\} \\ m_4 = \{l, c\} & m_5 = \{l, r\} & m_6 = \{c, r\}. \end{array}$$

Remember that there must be exactly two balls in the magic box. For this reason $\{l, c, r\}$ is not a local model describing a viewpoint of Mr. Pink.

Compatibility relations and model Following the definition given in Section 3, a generic compatibility pair for the magic box is a pair $\langle \mathbf{c}_B, \mathbf{c}_P \rangle$ where $\mathbf{c}_B$ is a set of models of the view of Mr. Blue and $\mathbf{c}_P$ is a set of models of the view of Mr. Pink. A model is a set of compatibility pairs.

In order to construct a model for the scenario described in Figure 3 (Section 2), we impose the following compatibility constraints:

$$\begin{array}{c} \text{if Mr. Blue sees a single blue ball} \\ \text{then Mr. Pink sees two pink balls} \end{array} \tag{1}$$

$$\begin{array}{c} \text{if Mr. Pink sees a single pink ball} \\ \text{then Mr. Blue sees two blue balls} \end{array} \tag{2}$$

$$\begin{array}{c} \text{Mr. Blue and Mr. Pink are able to construct} \\ \text{a complete description of their view} \end{array} \tag{3}$$

Notationally we use the following shorthand:

- $one(l, r)$ for $(l \vee r) \wedge \neg(l \wedge r)$;
- $one(l, c, r)$ for $(l \vee c \vee r) \wedge \neg(l \wedge r) \wedge \neg(l \wedge c) \wedge \neg(c \wedge r)$;
- $two(l, c, r)$ for $((l \wedge r) \vee (l \wedge c) \vee (c \wedge r)) \wedge \neg(l \wedge c \wedge r)$.

Constraints (1)-(3) are captured, at a formal level, by the following definition. A model $\mathbf{C}$ for the magic box is a compatibility relation such that, for all $\mathbf{c} \in \mathbf{C}$

$$\text{if } \mathbf{c}_B \text{ satisfies } one(l, r) \text{ then } \mathbf{c}_P \text{ satisfies } two(l, c, r) \tag{4}$$

$$\text{if } \mathbf{c}_P \text{ satisfies } one(l, c, r) \text{ then } \mathbf{c}_B \text{ satisfies } l \wedge r \tag{5}$$

$$|\mathbf{c}_B| = 1 \text{ and } |\mathbf{c}_P| = 1 \tag{6}$$

Let us explore in detail the relation between the informal compatibility constraints (1)-(3) and Equations (4)-(6). Equation (4) models constraint (1). In fact, if Mr. Blue sees a ball then this ball can be on the left or on the right and the formula $one(l, r)$ describes his view. Furthermore, in this case, Mr. Pink sees two balls in two of the three possible positions, and, therefore $two(l, c, r)$ represents his view. A similar explanation can be given for Equation (5), which models constraint (2). Equation (6) is more interesting. It says that $\mathbf{c}_B$ and $\mathbf{c}_P$ contain a single local model, i.e., the magic box model is a chain model. This intuitively means that both Mr. Blue and Mr. Pink see the box (from their point of view) and are able to construct a complete description of it. As a consequence of Equation (6), a model $\mathbf{C}$ for the magic box example in Figure 3 is a set of pairs $\langle \{m_B\}, \{m_P\} \rangle$ where m_B and m_P are local models of L_B and L_P, respectively. Each pair corresponds to a possible combination of the observers' partial views. The model $\mathbf{C}$ containing all and only the compatibility pairs depicted in Figure 3 is represented in Equation (7). All the models satisfying Equations (4)-(6) are subsets of this model.

$$\mathbf{C} = \left\{ \begin{array}{cc} \langle\{l\},\{l,c\}\rangle, & \langle\{l\},\{l,r\}\rangle, \\ \langle\{l\},\{c,r\}\rangle, & \langle\{r\},\{l,c\}\rangle, \\ \langle\{r\},\{l,r\}\rangle, & \langle\{r\},\{c,r\}\rangle, \\ \langle\{l,r\},\{l\}\rangle, & \langle\{l,r\},\{c\}\rangle, \\ \langle\{l,r\},\{r\}\rangle, & \langle\{l,r\},\{l,c\}\rangle, \\ \langle\{l,r\},\{l,r\}\rangle, & \langle\{l,r\},\{c,r\}\rangle \end{array} \right\} \tag{7}$$

As a final remark notice that linking local models inside a model may force us to eliminate some of them. Suppose that we restrict ourselves to consider local models for Mr. Blue which allow for exactly one ball. This leads to the definition of the two local models $\{l\}$ and $\{r\}$ for L_B depicted on the lefthand side in Figure 7, and of the six possible local models $\{l\}$, $\{c\}$, $\{r\}$, $\{l,c\}$, $\{l,r\}$, $\{c,r\}$ for L_P depicted on the righthand side in Figure 7. We know that if Mr. Blue sees a single ball, then Mr. Pink must see two balls.

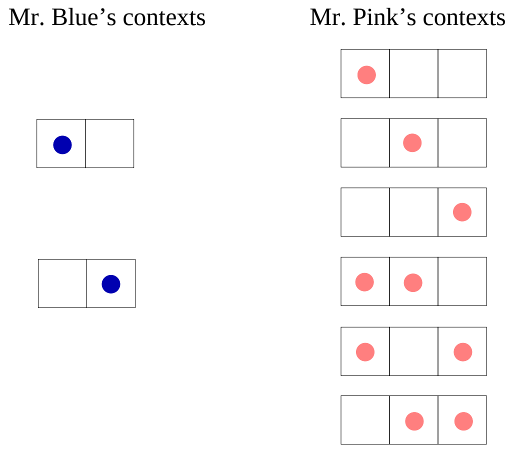

Figure 7: Mr. Blue sees exactly one ball: the local models.

As a consequence, the model for the situation in which Mr. Blue sees exactly one ball does not contain any pair, and corresponding local models for Mr. Pink, which represent that Mr. Pink sees a single ball. The resulting model is indeed the following:

$$\left\{ \begin{array}{cc} \langle\{l\},\{l,c\}\rangle, & \langle\{l\},\{l,r\}\rangle, \\ \langle\{l\},\{c,r\}\rangle, & \langle\{r\},\{l,c\}\rangle, \\ \langle\{r\},\{l,r\}\rangle, & \langle\{r\},\{c,r\}\rangle \end{array} \right\}$$

and is graphically represented in Figure 8.

4 Satisfiability and logical consequence

The definition of satisfiability of a formula of a language L_i in the model $\mathbf{C}$, is based on the satisfiability of the same formula in the local models of L_i. Formally, let $\models_{cl}$ be the satisfiability relation between local models and formulae of L_i. We call $\models_{cl}$ *local satisfiability*. Notationally, let us write $i{:}\phi$ to mean ϕ, where ϕ is a formula of L_i. We say that ϕ is an L_i-formula, and that $i{:}\phi$ is a formula or, also, a labelled L_i-formula. This notation and terminology allows us to keep track of the context we are talking about.

Let $\mathbf{C} = \{\mathbf{c}\}$ with $\mathbf{c} = \langle \mathbf{c}_0, \mathbf{c}_1, \ldots, \mathbf{c}_i, \ldots \rangle$ be a model and $i : \phi$ a formula. $\mathbf{C}$ *satisfies* $i : \phi$, in symbols $\mathbf{C} \models i{:}\phi$, if for all $\mathbf{c} \in \mathbf{C}$

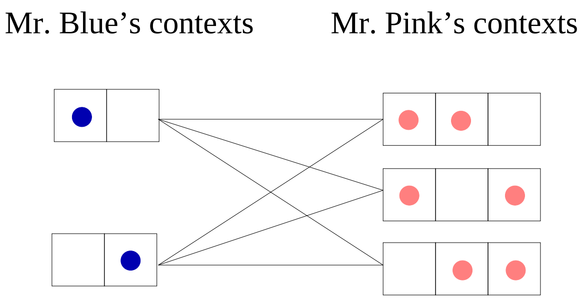

Figure 8: Mr. Blue sees exactly one ball: the model.

$$\mathbf{c}_i \models \phi$$

where $\mathbf{c}_i \models \phi$ if, for all $m \in \mathbf{c}_i$, $m \models_{cl} \phi$.

The intuition underlying the motion of satisfiability is that an L_i-formula is satisfied by a model $\mathbf{C}$ if all the local models in each $\mathbf{c}_i$ satisfy it.

Consider, for instance, the simple model

$$\mathbf{C}' = \{ \ \langle \{l\}, \{c, r\} \rangle, \quad \langle \{l, r\}, \{l, c\} \rangle \ \} \tag{8}$$

containing only the two compatibility pairs depicted in Figure 9. According to the definition of satisfiability $\mathbf{C}'$ satisfies the formula $B\!:l$, meaning that Mr. Blue sees a ball in the left position. This is because the two local models $\{l\}$ and $\{l, r\}$ for L_B contained in $\mathbf{C}'$ both satisfy the formula l. On the contrary, $\mathbf{C}'$ does not satisfy $B\!:r$, meaning that Mr. Blue sees a ball in the right position. This is because there is a local model for Mr. Blue, namely $\{l\}$, which does not satisfy the formula r.

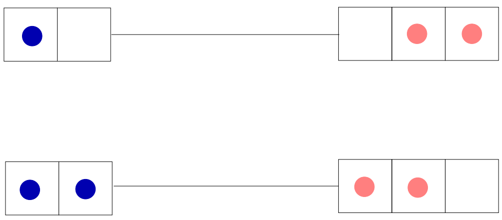

Figure 9: Mr. Blue sees a ball on the left.

The notions of satisfiability of a set of formulae and of validity are the obvious ones. A model $\mathbf{C}$ satisfies a set of formulae Γ, in symbols $\mathbf{C} \models \Gamma$, if $\mathbf{C}$ satisfies every formula $i\!:\phi$ in Γ. A formula $i\!:\phi$ is *valid*, in symbols $\models i\!:\phi$, if all models satisfy $i\!:\phi$.

An interesting notion is the one of *logical consequence* which must take into account the fact that assumptions and conclusion may belong to distinct languages. Given a set of labelled formulae Γ, Γ_j denotes the set of formulae $\{\gamma \mid j\!:\gamma \in \Gamma\}$. A formula $i\!:\phi$ is a logical consequence of a set of formulae Γ w.r.t. a model $\mathbf{C}$, in symbols $\Gamma \models_{\mathbf{C}} i\!:\phi$, if every sequence $\mathbf{c} \in \mathbf{C}$ satisfies:

$$\forall j \in I, j \neq i, \ \mathbf{c}_j \models \Gamma_j \Longrightarrow (\forall m \in \mathbf{c}_i, \ m \models_{cl} \Gamma_i \Longrightarrow m \models_{cl} \phi) \tag{9}$$

Equation (9) looks slightly complicated. Let us illustrate it with the help of an example. Consider the model of the magic box informally depicted in Figure 3 and formally represented by Equation (7). We want to verify that in this model

(10) if Mr. Blue sees a ball on the left and no ball on the right, and Mr. Pink doesn't see any ball in the center, then Mr. Pink sees a ball on the left and a ball on the right.

Formally, the sentence (10) can be rewritten as

$$B{:}\,l \wedge \neg r, P{:}\,\neg c \models_{\mathbf{C}} P{:}\,l \wedge r$$

The set of assumption Γ contains the facts that "Mr. Blue sees a ball on the left and no ball on the right" and "Mr. Pink doesn't see any ball in the center". Formally, $\Gamma = \{B{:}\,l \wedge \neg r, P{:}\,\neg c\}$. The first step is to isolate the set of assumptions which are made in a context different from the context of Mr. Pink. That is $B{:}\,l \wedge \neg r$. Then we restrict ourselves to considering all the compatibility pairs whose local models satisfy the formula $B{:}\,l \wedge \neg r$, and throw away all the others. The remaining compatibility pairs are

$$\langle\{l\}, \{l, c\}\rangle$$
$$\langle\{l\}, \{l, r\}\rangle$$
$$\langle\{l\}, \{c, r\}\rangle$$

and are depicted in Figure 10. Consider now the local models of Mr. Pink in the remaining sequences. We

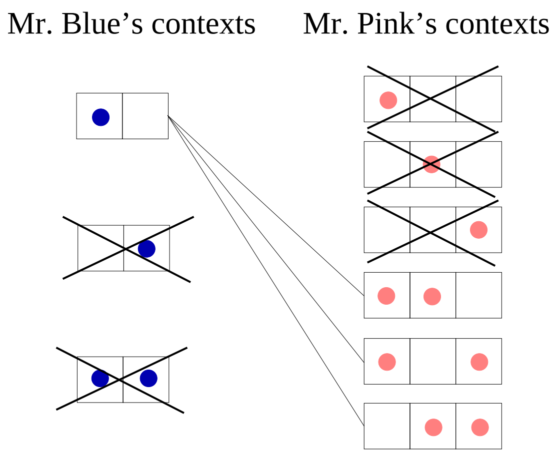

Figure 10: Selecting compatibility sequences.

have to identify all the local models of Mr. Pink in the remaining pairs such that there is no ball in the center. Formally, we have to identify all the local models of Mr. Pink satisfying $P{:}\,\neg c$. The only local model satisfying that Mr. Pink doesn't see any ball in the center is

$$\{l, r\}$$

and is depicted in Figure 11. The last step is to check whether the remaining local models of Mr. Pink

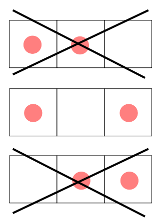

Figure 11: Selecting local models.

represent the fact that Mr. Pink sees a ball on the left and a ball on the right. It is easy to see that the only remaining local model in Figure 11 satisfies this property. Therefore the model depicted in Figure 3 and formally defined in Equation (7) satisfies the sentence (10).

The extension of the notion of logical consequence to a class of models is the usual one. A formula $i\colon\phi$ is a logical consequence of a set of formulae Γ w.r.t. a class of models M, in symbols $\Gamma \models_{\mathsf{M}} i\colon\phi$, if $i\colon\phi$ is a logical consequence of Γ w.r.t. all the models in M. Finally, a formula $i\colon\phi$ is a *logical consequence* of Γ, in symbols $\Gamma \models i\colon\phi$, if $i\colon\phi$ is a logical consequence of Γ w.r.t. all models $\mathbf{C}$.

5 Contexts, locality and compatibility

Having formally defined the logical framework, the question now is: where are contexts in this picture? How does Local Models Semantics relate to contextual reasoning? We already suggested part of the answer to this question by illustrating the main notions of model and satisfiability using the magic box example. In this section we answer these questions in more detail by illustrating how the notion of context can be formally introduced in the framework of Local Models Semantics. We then examine how Local Models Semantics formally captures the notions of locality and compatibility.

Given a model $\mathbf{C} = \{\langle \mathbf{c}_0, \mathbf{c}_1, \ldots, \mathbf{c}_i, \ldots \rangle\}$ we formally define a *context* to be any $\mathbf{c}_i$, namely the set of local models $m \in \overline{M}_i$ allowed by $\mathbf{C}$ within any particular compatibility sequence. For instance, the contexts for Mr. Blue allowed by the model $\mathbf{C}'$ defined in Equation (8) are $\{l\}$ and $\{l, r\}$.

The intuition underlying the definition of context is that a context consists of that set of models which capture exactly those facts which are locally true, given also the constraints posed by the local models of other contexts in the same compatibility sequence. This notion of context is the semantic formalisation of the notion of context intuitively introduced in Principle 1 in Section 1.

An interesting property of this definition is that contexts are formalised as partial objects, as explicitly required in, e.g., [14; 19]. This is due to the fact that context is defined as a set of models instead of a single model. In order to illustrate the advantage of having contexts as partial objects consider the slightly modified magic box scenario depicted in Figure 12, where Mr. Pink is able to see only one box sector and knows that there are two sectors behind the wall. In this scenario Mr. Pink is able to distinguish only two

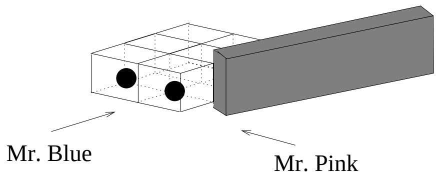

Figure 12: A partially hidden magic box.

situations: there is a ball on the left, and there is no ball on the left. The fact that Mr. Pink is uncommitted to whether there is a ball in a sector behind the wall is formalised by having the sentences "there is a ball on the right" and "there is a ball in the center" true in some local models representing the view of Mr. Pink and false in others. In the resulting context, describing the viewpoint of Mr. Pink, "there is a ball on the right" and "there is a ball in the center" will be neither true or false because there will be models in $\mathbf{c}_P$ where these sentences are false and others where the same sentences are true. Formally, the model for the scenario depicted in Figure 12 is defined as follows

$$\mathbf{C}^* = \left\{ \begin{array}{c} \langle \{l\}, \{\{c\}, \{r\}, \{c, r\}\} \rangle, \\ \langle \{r\}, \{\{c\}, \{r\}, \{c, r\}\} \rangle, \\ \langle \{l, r\}, \{\{c\}, \{r\}, \{c, r\}\} \rangle, \\ \langle \{l\}, \{\{l\}, \{l, c\}, \{l, r\}\} \rangle, \\ \langle \{r\}, \{\{l\}, \{l, c\}, \{l, r\}\} \rangle, \\ \langle \{l, r\}, \{\{l\}, \{l, c\}, \{l, r\}\} \rangle \end{array} \right\} \tag{11}$$

and is graphically represented in Figure 13. It is easy to see that the two contexts for Mr. Pink allowed by

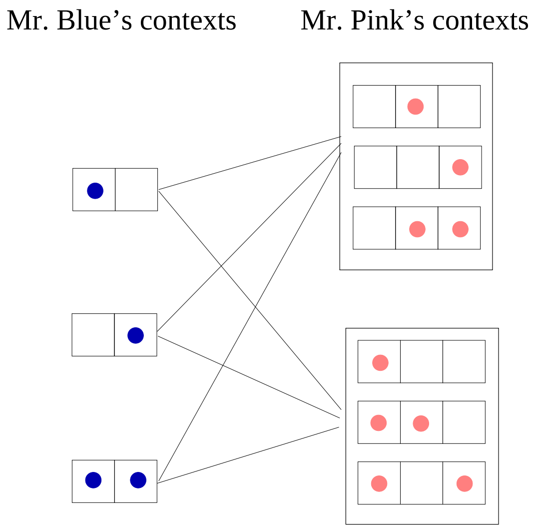

Figure 13: Model for the scenario of Figure 12.

the model $\mathbf{C}^*$ are $\{\{c\},\{r\},\{c,r\}\}$ and $\{\{l\},\{l,c\},\{l,r\}\}$. In these contexts the formulae r and c are neither true or false. Consider, for instance, the context $\{\{c\},\{r\},\{c,r\}\}$ and the formula r. r is neither true or false in $\{\{c\},\{r\},\{c,r\}\}$ because there is a local model $\{c\}$ where r is false and another local model $\{r\}$ where r is true.

Given the above notion of context, we can now better illustrate the intuitions underlying the notions of compatibility sequence, compatibility relation, and model. A context is a partial description of the world. A compatibility sequence contains as many contexts as needed, one for each partial description of the world. Thus, in the magic box scenario we have compatibility sequences of length two, containing a context for the view of Mr. Blue and a context for the view of Mr. Pink. In the more general scenario involving n observers, we have to consider sequences of length n.

An interesting set of compatibility sequences is the one composed by chains introduced at the end of Section 3. Remember that a chain is a compatibility sequence in which all the contexts are singleton sets. In this case, all the contexts are complete objects in the sense that each context, being a single model, assigns a truth value to all sentences in its language. In other words, a context which is a singleton set models the situation where a partial description of the world assigns a truth value to all the propositions it is able to express in its local (and limited) language. This is the case in Figures 1, 2, and 3. Here, Mr. Blue and Mr. Pink have partial views of the world. However, within their partial views, they are able to "see everything". On the contrary, this is not the case in Figures 12 and 13. Here, Mr. Blue is still able to "see everything" within its partial views, while Mr. Pink is not.

Local Models Semantics completely embraces the principle of Locality. We can easily say that everything is local. First of all, the languages are local to the contexts. Second, the languages are interpreted in local structures (or local models). This reflects the fact that contexts can have their own, generally different, domains of interpretation, sets of relations, and sets of functions. Third, the notion of satisfiability is local: the satisfiability of a (labelled) formula is given in terms of the local satisfiability of the formula with respect to its context.

Because of compatibility sequences, contexts mutually influence themselves. Compatibility has the structural effect of changing the set of local models defining each context. It forces local models to agree up to a certain extent. A typical example is the one depicted in Figure 8, where the fact that Mr. Blue sees exactly a ball forces us to throw away all the pairs, and corresponding local models for Mr. Pink, which allow for zero balls.

6 From contexts to the world

In learning about our approach to the formalisation of the magic box example, the reader might object that the most straightforward formalisation of this example would be a direct axiomatisation of the box as a two-dimensional grid. The contexts representing the views of Mr. Blue and Mr. Pink could then easily

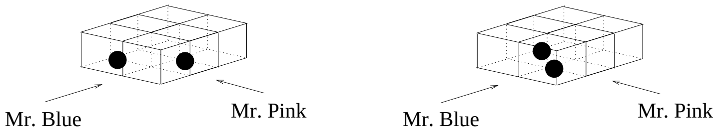

Figure 14: Indistinguishable situations.

be constructed by projecting the grid in two one-dimensional views. Locality and compatibility would be guaranteed by construction. However this approach is based on the hypothesis that we have a complete description of the world (the box in this case), and that we can use it to build views of the world itself. This is not always the case. Quite often we have only partial views and it is possible that we are not able to reconstruct the complete description of the world starting from the partial views, but only a partial or approximate description of it. As an example, consider the situations depicted in Figure 14. These two different situations cannot be distinguished by the two observers. That is, even assuming the existence of a third agent who knows the actual form of the box, (s)he is not able to identify which situation, among the ones depicted in Figure 14, is the current one, knowing only what Mr. Blue and Mr. Pink see. In fact, the unique pair of compatible contexts associated to the two different situations in Figure 14 is the one depicted in Figure 15.

Figure 15: Compatible contexts in the scenario of Figure 14.

The capability of dealing with situations where we may or may not have a complete description of the world is quite important in several application domains. Among the most important is the development and integration of data or knowledge bases. In a relational, possibly distributed, data base there is what is assumed to be a complete description of the world, and views are built by filtering out, and appropriately merging together, part of the available information. On the other hand, a federation of heterogeneous data or knowledge bases, possibly developed independently, can be seen as a set of views of an ideal data base which is often impossible or very complex to reconstruct completely.

An exhaustive investigation on the relation between partial views and a complete description of the world is out of the scope of this chapter. Our aim here is to highlight the problem and suggest how Local Models Semantics is able to deal with situations where we may or may not have a complete description of the world in (simple) scenarios from the magic box example. In order to do that, consider the following scenario. The box is the same as the one depicted in Figure 1, but this time the balls have to be placed in the same column (i.e., there cannot be balls on a diagonal line). Figure 16 shows all the possible configurations allowed in this scenario from a top view of the box.

It is very easy to show that in this case the observers can distinguish between all the possible situations. Figure 17 graphically describes the compatibility pairs involving the three different possible situations for Mr. Blue and the six different possible situations for Mr. Pink.

The graphical model depicted in Figure 17 doesn't look very different from the one depicted in Figure 3. So, why in this case the observers are able to distinguish between all the possible situations? Because in

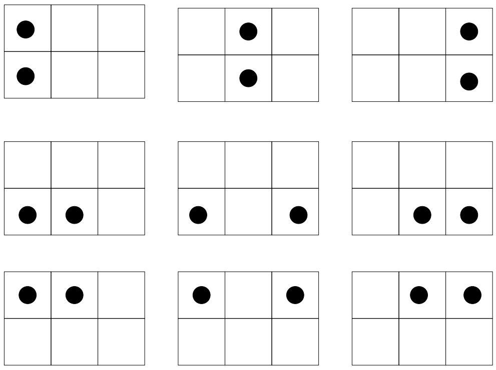

Figure 16: A new magic box.

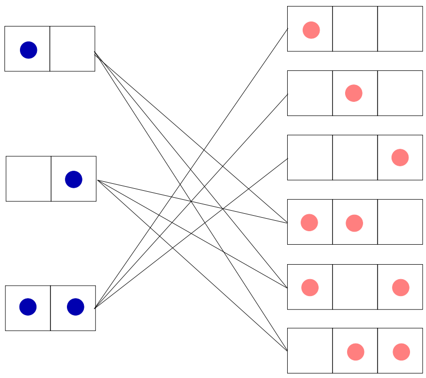

Figure 17: Compatible contexts in the scenario of Figure 16.

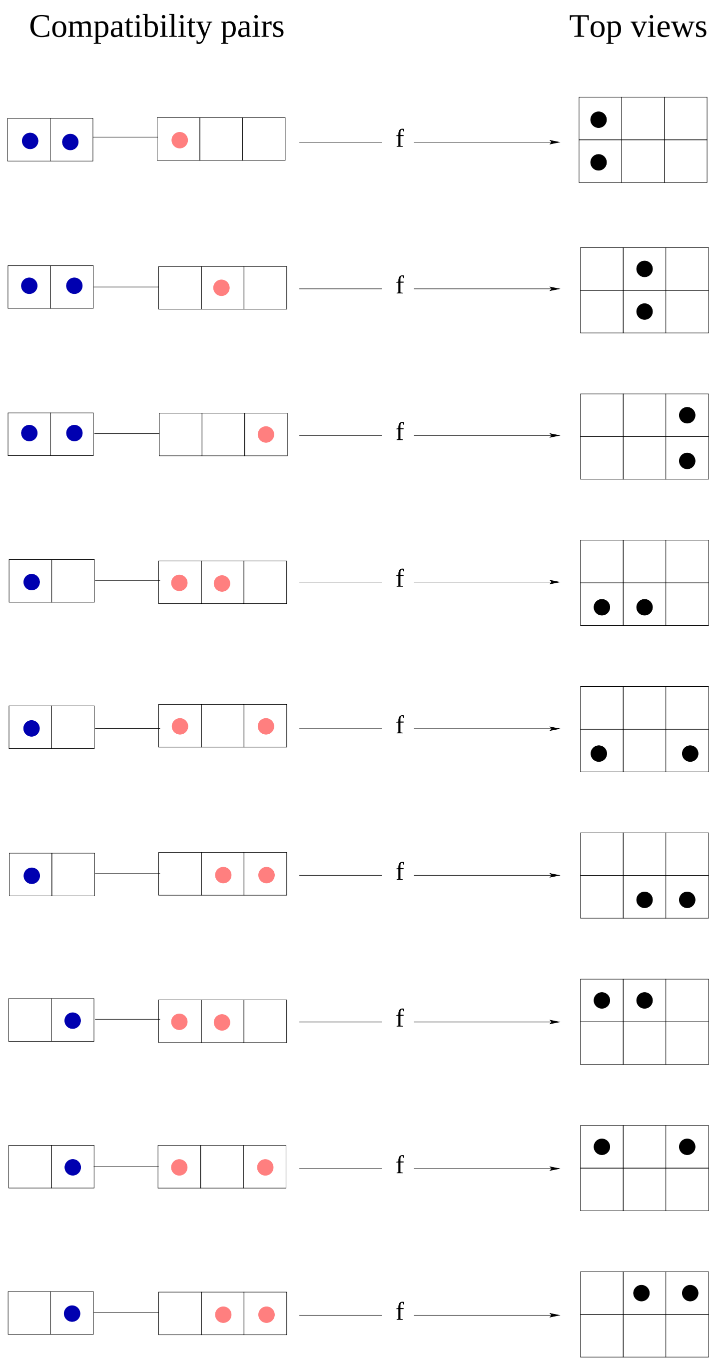

Figure 18: One-to-one correspondence.

this case it is possible to find a precise correspondence between the compatibility pairs in Figure 17 and the complete description of the box provided by the top views in Figure 16. More formally, it is possible to find a bijective[4] function f from the set of compatibility pairs $\mathbf{C}$, graphically defined in Figure 17, to the set of models graphically defined in Figure 16. This function enables a one-to-one correspondence between every compatibility pair in Figure 17 and one of the possible descriptions of the box, in Figure 16. Figure 18 provides a graphical description of f.

Let $\mathbf{C}$ be a compatibility relation and M a set of models intuitively representing a complete description of the world. We believe that the capability of defining a bijective function f from $\mathbf{C}$ to M is a necessary condition for stating that $\mathbf{C}$ enables the reconstruction of a complete description of the world. Is this condition also a sufficient one? Due to the infinite varieties of relations existing between different views of the world we are not able to give a definite answer in this chapter. Nonetheless, one-to-one functions can provide a preliminary mechanism for controlling whether a certain model $\mathbf{C}$ provides a description of different views of the world which enables the reconstruction of a complete description of the world.

7 Conclusion

In this chapter we have explained a new semantics, called Local Models Semantics, which was recently proposed as a foundation to reasoning with context. Local Models Semantics formalises the two general principles underlying contextual reasoning, namely the principle of locality and the principle of compatibility. We have also shown how Local Models Semantics can be used to model a characteristic example of reasoning with viewpoints: the magic box example.

References

[1] V. Akman and M. Surav. Steps toward Formalizing Context. *AI Magazine*, pages 55–72, FALL 1996.

[2] G. Attardi and M. Simi. A formalisation of viewpoints. *Fundamenta Informaticae*, 23(2–4):149–174, 1995.

[3] M. Benerecetti, P. Bouquet, and C. Ghidini. Formalizing Belief Reports – The Approach and a Case Study. In F. Giunchiglia, editor, *Proceedings AIMSA'98, 8th International Conference on Artificial Intelligence, Methodology, Systems, and Applications*, volume 1480 of *LNAI*, pages 62–75, Sozopol, Bulgaria, September, 21–23 1998. Springer Verlag.

[4] 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.

[5] P. Bouquet and F. Giunchiglia. Reasoning about Theory Adequacy. A New Solution to the Qualification Problem. *Fundamenta Informaticae*, 23(2–4):247–262, June,July,August 1995.

[6] A. Cimatti and L. Serafini. Multi-Agent Reasoning with Belief Contexts: the Approach and a Case Study. In M. Wooldridge and N. R. Jennings, editors, *Intelligent Agents: Proceedings of 1994 Workshop on Agent Theories, Architectures, and Languages*, number 890 in Lecture Notes in Computer Science, pages 71–85. Springer Verlag, 1995.

[7] A. Farquhar, A. Dappert, R. Fikes, and W. Pratt. Integrating Information Sources Using Context Logic. In *Proceedings of AAAI Spring Symposium on Information Gathering from Distributed Heterogeneous Environments*, 1995.

[4] Formally, a function f from a set A to a set B is *injective* if each element of A maps onto a different element of B. A function f from set A onto B is called *surjective* (or 'onto') if every member of B is the image of at least one member of A. A function f is *bijective* if it is both injective and surjective.

[8] M. Fisher and C. Ghidini. Programming Resource-Bounded Deliberative Agents. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI'99)*, pages 200–206. Morgan Kaufmann Publ., Inc, 1999.

[9] C. Ghidini. Modelling (Un)Bounded Beliefs. In P. Bouquet, L. Serafini, P. Brezillon, M. Benerecetti, and F. Castellani, editors, *Modelling and Using Context – Proceedings of the 2nd International and Interdisciplinary Conference, Context'99*, volume 1688 of *Lecture Notes in Artificial Intelligence*, pages 145–158. Springer Verlag - Heidelberg, 1999.

[10] C. Ghidini and F. Giunchiglia. Local models semantics, or contextual reasoning = locality + compatibility. *Artificial Intelligence*, 127(2):221–259, April 2001.

[11] C. Ghidini and L. Serafini. Information Integration for Electronic Commerce. In C. Sierra and P. Noriega, editors, *Proceedings of the Workshop on Agent Mediated Electronic Trading (AMET'98)*, volume 1571 of *LNAI*, pages 189–206, Minneapolis, USA, May, 10 1998. Springer Verlag.

[12] C. Ghidini and L. Serafini. Model Theoretic Semantics for Information Integration. In F. Giunchiglia, editor, *Proceedings AIMSA'98, 8th International Conference on Artificial Intelligence, Methodology, Systems, and Applications*, volume 1480 of *LNAI*, pages 267–280, Sozopol, Bulgaria, September, 21–23 1998. Springer Verlag.

[13] E. Giunchiglia and F. Giunchiglia. Ideal and Real Belief about Belief. In *Practical Reasoning, International Conference on Formal and Applied Practical Reasoning, FAPR'96*, number 1085 in Lecture Notes in Artificial Intelligence, pages 261–275. Springer Verlag, 1996.

[14] F. Giunchiglia. Contextual reasoning. *Epistemologia, special issue on I Linguaggi e le Macchine*, XVI:345–364, 1993. Short version in Proceedings IJCAI'93 Workshop on Using Knowledge in its Context, Chambery, France, 1993, pp. 39–49.

[15] F. Giunchiglia and L. Serafini. Multilanguage hierarchical logics, or: how we can do without modal logics. *Artificial Intelligence*, 65(1):29–70, 1994.

[16] F. Giunchiglia, L. Serafini, E. Giunchiglia, and M. Frixione. Non-Omniscient Belief as Context-Based Reasoning. In *Proc. of the 13th International Joint Conference on Artificial Intelligence*, pages 548–554, Chambery, France, 1993.

[17] R.V. Guha. *Contexts: a Formalization and some Applications.* PhD thesis, Department of Computer Science, Stanford University, 1991. Also report No. STAN-CS-91-1399-Thesis.

[18] D. B. Lenat. Cyc: A large-scale investment in knowledge infrastructure. *Communications of the ACM*, 38(11):33–38, 1995.

[19] J. McCarthy. Generality in Artificial Intelligence. *Communications of ACM*, 30(12):1030–1035, 1987. Also in V. Lifschitz (ed.), *Formalizing common sense: papers by John McCarthy*, Ablex Publ., 1990, pp. 226–236.

[20] J. McCarthy. Notes on Formalizing Context. In *Proc. of the 13th International Joint Conference on Artificial Intelligence*, pages 555–560, Chambery, France, 1993.

[21] J. Mylopoulos and R. Motschnig-Pitrik. Partitioning Information Bases with Contexts. In *Third International Conference on Cooperative Information Systems*, Vienna, 1995.

[22] P. Noriega and C. Sierra. Towards Layered Dialogical Agents. In J. P. Müller, M. J. Wooldridge, and N. R. Jennings, editors, *Proceedings of the Third International Workshop on Agent Theories, Architectures, and Languages (ATAL-96), ECAI-96*, volume 1193 of *LNAI*, pages 173–189, Budapest, Hungary, 1996.

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

[24] M. Theodorakis, A. Analyti, P. Constantopoulos, and N. Spyratos. Context in Information Bases. In *Proceedings of the 3rd International Conference on Cooperative Information Systems (CoopIS'98)*, pages 260–270, New York City, NY, USA, August 1998. IEEE Computer Society.