

CSLI BOOK TITLE

Paolo Bouquet and Luciano Serafini (eds.)

CENTER FOR THE STUDY
OF LANGUAGE
AND INFORMATION

Contents

1	What is Local Models Semantics?	1
	CHIARA GHIDINI, FAUSTO GIUNCHIGLIA	
2	On the dimensions of context dependence	25
	MASSIMO BENERECETTI, PAOLO BOUQUET, CHIARA GHIDINI	

November 8, 2002

1

What is Local Models Semantics?

CHIARA GHIDINI, FAUSTO GIUNCHIGLIA

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.

1.1 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 (1993) and John McCarthy (1993) in the late 80's, when context was introduced as an important means for formalising certain forms of reasoning.

According to Giunchiglia (1993), 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

¹The interested reader may refer to Akman and Surav (1996), Ghidini and Giunchiglia (2001), Giunchiglia (1993) for an accurate discussion on this topic.

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) Giunchiglia (1993).

In McCarthy (1993), 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. McCarthy (1987)

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 Farquhar et al. (1995), Mylopoulos and Motschnig-Pitrik (1995), Ghidini and Serafini (1998b), Theodorakis et al. (1998). The largest common-sense knowledge-base, CYC Lenat (1995), contains an explicit notion of context Guha (1991). 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. See for instance Giunchiglia et al. (1993), Giunchiglia and Giunchiglia (1996), Benerecetti et al. (1998a), Fisher and Ghidini (1999), Ghidini (1999), Giunchiglia and Serafini (1994). In Attardi and Simi (1995) contexts are introduced in the formalisation of reasoning with viewpoints. Bouquet and Giunchiglia (1995) addresses the problem of formalising context-based common-sense reasoning. Finally, Benerecetti et al. (1998b), Cimatti and Serafini (1995), Parsons et al. (1998), Noriega and Sierra (1996), Ghidini and Serafini (1998a) 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 Ghidini and Giunchiglia (2001) 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 Ghidini and Giunchiglia (2001)

The chapter is organised as follows. The core definitions are given in Sections 1.3 and Section 1.4. In Section 1.5 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 1.6 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 1.2 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 Ghidini and Giunchiglia (2001).

1.2 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

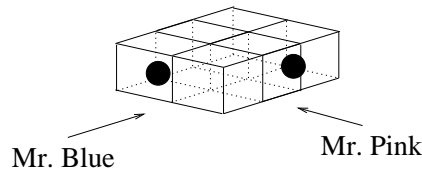


FIGURE 1 The magic box.

sectors, each sector possibly containing a ball. There must be exactly 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 view-

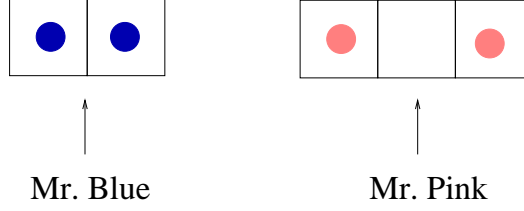


FIGURE 2 The contexts of Mr. Blue and Mr. Pink.

points 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, 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 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”.

1.3 Local models and model

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

Mr. Blue's contexts Mr. Pink's contexts

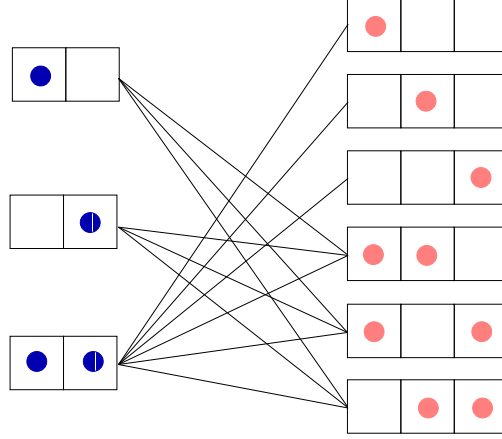


FIGURE 3 Compatible contexts of Mr. Blue and Mr. Pink.

1.3.1 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

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

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

²Formally, let $\prod_{i \in I} 2^{\overline{M}_i}$ be the Cartesian product of the collection $\{2^{\overline{M}_i} : i \in I\}$.

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. $\langle \emptyset, \emptyset, \dots, \emptyset, \dots \rangle \notin \mathbf{C}$.

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.

1.3.2 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 \perp , and closed under implication; L_P is formally defined as the smallest set containing P_2 , the symbol for falsity \perp and closed under implication. In this chapter we use the standard abbreviations from propositional logic, such as $\neg\phi$ for $\phi \supset \perp$, $\phi \vee \psi$ for $\neg\phi \supset \psi$, $\phi \wedge \psi$ for $\neg(\neg\phi \vee \neg\psi)$, \top for $\perp \supset \perp$.

The compatibility relation \mathbf{C} is a relation of type $\mathbf{C} \subseteq \prod_{i \in I} 2^{\overline{M}_i}$

³Figures 4, 5, and 6 first appeared in Ghidini and Giunchiglia (2001).

WHAT IS LOCAL MODELS SEMANTICS? / 7



FIGURE 4 Languages: L_1 , L_2 , and L_3 .

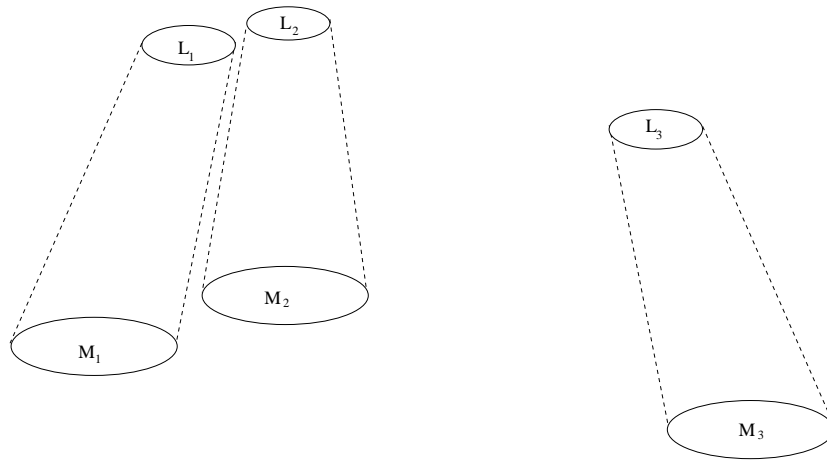


FIGURE 5 Local models for L_1 , L_2 , and L_3 .

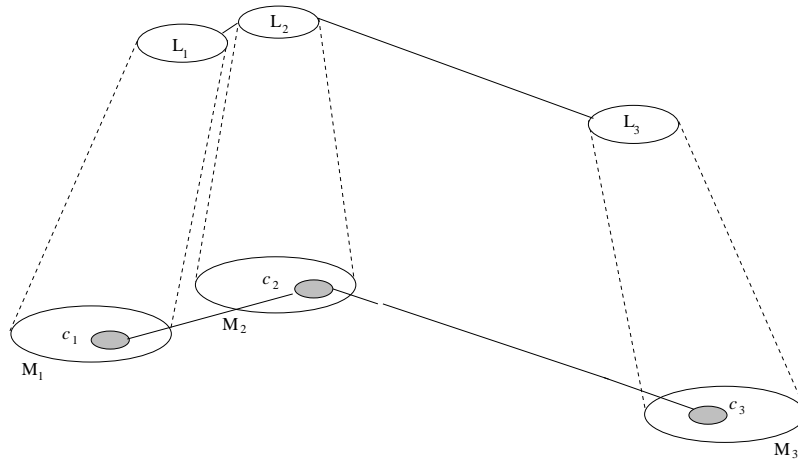


FIGURE 6 Model for $\{L_1, L_2, L_3\}$.

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:

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

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 1.3, 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 1.2), we impose the following compatibility constraints:

$$\begin{array}{ll} \text{if Mr. Blue sees a single blue ball} & \\ \text{then Mr. Pink sees two pink balls} & (1.1) \end{array}$$

$$\begin{array}{ll} \text{if Mr. Pink sees a single pink ball} & \\ \text{then Mr. Blue sees two blue balls} & (1.2) \end{array}$$

$$\begin{array}{ll} \text{Mr. Blue and Mr. Pink are able to construct} & \\ \text{a complete description of their view} & (1.3) \end{array}$$

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 (1.1)-(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 $\mathbf{c} \in \mathbf{C}$

$$\text{if } \mathbf{c}_B \text{ satisfies } one(l, r) \text{ then } \mathbf{c}_P \text{ satisfies } two(l, c, r) \quad (1.4)$$

$$\text{if } \mathbf{c}_P \text{ satisfies } one(l, c, r) \text{ then } \mathbf{c}_B \text{ satisfies } l \wedge r \quad (1.5)$$

$$|\mathbf{c}_B| = 1 \text{ and } |\mathbf{c}_P| = 1 \quad (1.6)$$

Let us explore in detail the relation between the informal compatibility constraints (1.1)-(1.3) and Equations (1.4)-(1.6). Equation (1.4) models constraint (1.1). 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 (1.5), which models constraint (1.2). Equation (1.6) 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 (1.6), 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 (1.7). All the models satisfying Equations (1.4)-(1.6) are subsets of this model.

$$\mathbf{C} = \left\{ \begin{array}{ll} \langle \{l\}, \{l, c\} \rangle, & \langle \{l\}, \{l, r\} \rangle, \\ \langle \{l\}, \{c, r\} \rangle, & \langle \{r\}, \{l, c\} \rangle, \\ \langle \{r\}, \{l, r\} \rangle, & \langle \{r\}, \{c, r\} \rangle, \\ \langle \{l, r\}, \{l\} \rangle, & \langle \{l, r\}, \{c\} \rangle, \\ \langle \{l, r\}, \{r\} \rangle, & \langle \{l, r\}, \{l, c\} \rangle, \\ \langle \{l, r\}, \{l, r\} \rangle, & \langle \{l, r\}, \{c, r\} \rangle \end{array} \right\} \quad (1.7)$$

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. 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

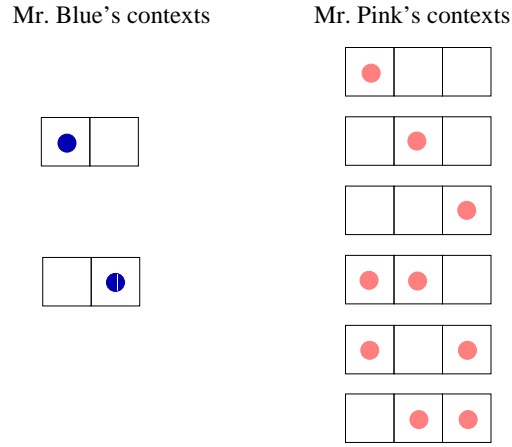


FIGURE 7 Mr. Blue sees exactly one ball: the local models.

Mr. Pink sees a single ball. The resulting model is indeed the following:

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

and is graphically represented in Figure 8.

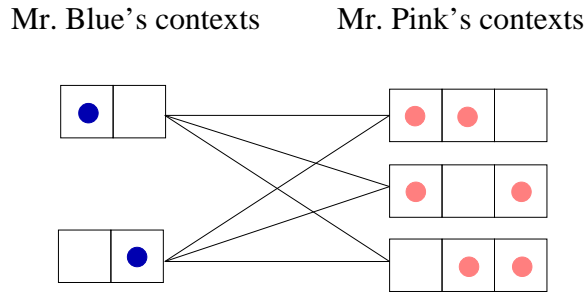


FIGURE 8 Mr. Blue sees exactly one ball: the model.

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

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

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

The intuition underlying the notion 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

$$\mathbf{C}' = \{ \langle \{l\}, \{c, r\} \rangle, \langle \{l, r\}, \{l, c\} \rangle \} \quad (1.8)$$

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 .

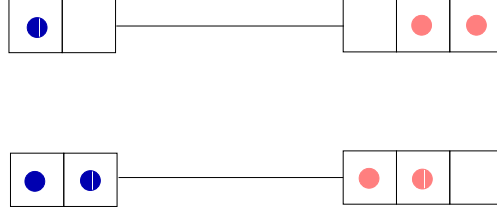


FIGURE 9 Mr. Blue sees a ball on the left.

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:

$$\forall j \in I, j \neq i, \mathbf{c}_j \models \Gamma_j \implies (\forall m \in \mathbf{c}_i, m \models_{cl} \Gamma_i \implies m \models_{cl} \phi) \quad (1.9)$$

Equation (1.9) 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 (1.7). 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

$$\begin{aligned} &\langle \{l\}, \{l, c\} \rangle \\ &\langle \{l\}, \{l, r\} \rangle \\ &\langle \{l\}, \{c, r\} \rangle \end{aligned}$$

and are depicted in Figure 10. Consider now the local models of Mr. Pink in the remaining sequences. We 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 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 (1.7) satisfies the sentence (10).

WHAT IS LOCAL MODELS SEMANTICS? / 13

Mr. Blue's contexts Mr. Pink's contexts

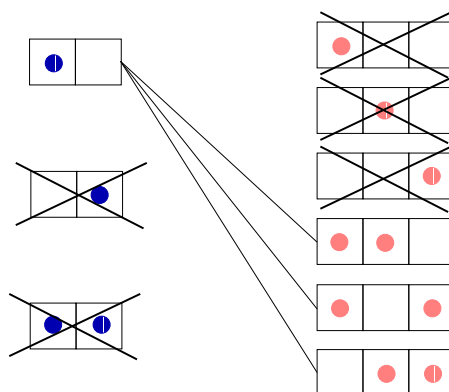


FIGURE 10 Selecting compatibility sequences.

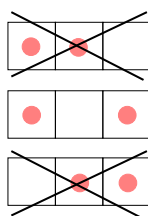


FIGURE 11 Selecting local models.

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

1.5 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 $\mathbf{C} = \{\langle \mathbf{c}_0, \mathbf{c}_1, \dots, \mathbf{c}_i, \dots \rangle\}$ 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 (1.8) 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 1.1.

An interesting property of this definition is that contexts are formalised as partial objects, as explicitly required in, e.g., Giunchiglia (1993), McCarthy (1987). 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 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

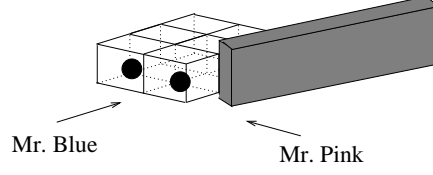


FIGURE 12 A partially hidden magic box.

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

$$\mathbf{C}^* = \left\{ \begin{array}{l} \langle \{l\}, \{\{c\}, \{r\}, \{c, r\}\}\rangle, \\ \langle \{r\}, \{\{c\}, \{r\}, \{c, r\}\}\rangle, \\ \langle \{l, r\}, \{\{c\}, \{r\}, \{c, r\}\}\rangle, \\ \langle \{l\}, \{\{l\}, \{l, c\}, \{l, r\}\}\rangle, \\ \langle \{r\}, \{\{l\}, \{l, c\}, \{l, r\}\}\rangle, \\ \langle \{l, r\}, \{\{l\}, \{l, c\}, \{l, r\}\}\rangle \end{array} \right\} \quad (1.11)$$

and is graphically represented in Figure 13. It is easy to see that the two contexts for Mr. Pink allowed by 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 1.3. Remember that a chain is a compatibility sequence in which all the contexts are singleton sets. In

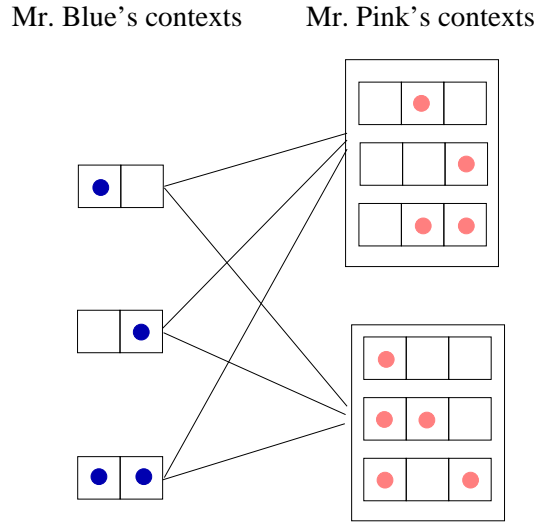


FIGURE 13 Model for the scenario of Figure 12.

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.

1.6 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

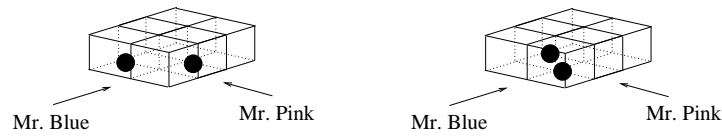


FIGURE 14 Indistinguishable situations.

and Mr. Pink could then easily 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.

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,



FIGURE 15 Compatible contexts in the scenario of Figure 14.

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.

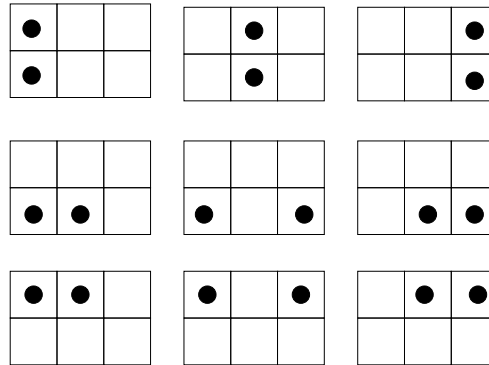


FIGURE 16 A new magic box.

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 differ-

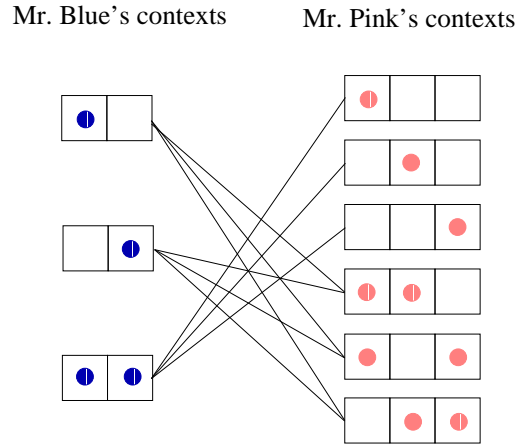


FIGURE 17 Compatible contexts in the scenario of Figure 16.

ent 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 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 bijective⁴ 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

⁴Formally, 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.

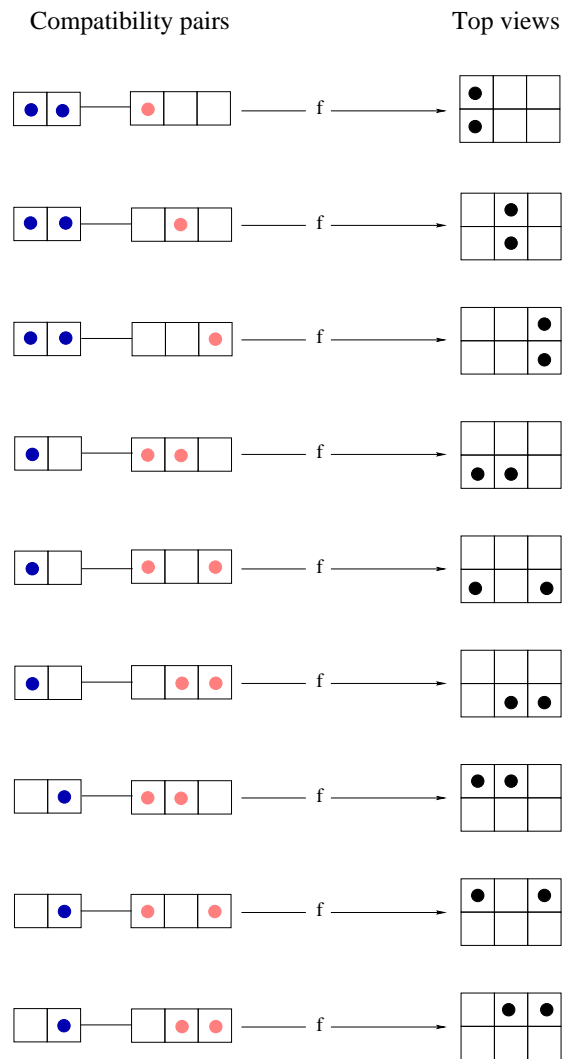


FIGURE 18 One-to-one correspondence.

of different views of the world which enables the reconstruction of a complete description of the world.

1.7 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. Chapter 2 (Section 2.5) contains a brief description of two additional, and very different, areas where Local Models Semantics has been successfully applied: the modelling of intentional context, and belief context in particular, and the representation of semantic heterogeneity issues in information integration.

References

- Akman, V. and M. Surav. 1996. Steps toward Formalizing Context. *AI Magazine* pages 55–72.
- Attardi, G. and M. Simi. 1995. A formalisation of viewpoints. *Fundamenta Informaticae* 23(2–4):149–174.
- Benerecetti, M., P. Bouquet, and C. Ghidini. 1998a. Formalizing Belief Reports – The Approach and a Case Study. In F. Giunchiglia, ed., *Proceedings AIMS’98, 8th International Conference on Artificial Intelligence, Methodology, Systems, and Applications*, vol. 1480 of *LNAI*, pages 62–75. Sozopol, Bulgaria: Springer Verlag.
- Benerecetti, M., F. Giunchiglia, and L. Serafini. 1998b. Model Checking Multiagent Systems. *Journal of Logic and Computation, Special Issue on Computational & Logical Aspects of Multi-Agent Systems* 8(3):401–423.
- Bouquet, P. and F. Giunchiglia. 1995. Reasoning about Theory Adequacy. A New Solution to the Qualification Problem. *Fundamenta Informaticae* 23(2–4):247–262.
- Cimatti, A. and L. Serafini. 1995. Multi-Agent Reasoning with Belief Contexts: the Approach and a Case Study. In M. Wooldridge and N. R. Jennings, eds., *Intelligent Agents: Proceedings of 1994 Workshop on Agent Theories, Architectures, and Languages*, no. 890 in *Lecture Notes in Computer Science*, pages 71–85. Springer Verlag.
- Farquhar, A., A. Dappert, R. Fikes, and W. Pratt. 1995. Integrating Information Sources Using Context Logic. In *Proceedings of AAAI Spring Symposium on Information Gathering from Distributed Heterogeneous Environments*.

- Fisher, M. and C. Ghidini. 1999. 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.
- Ghidini, C. 1999. Modelling (Un)Bounded Beliefs. In P. Bouquet, L. Serafini, P. Brezillon, M. Benerecetti, and F. Castellani, eds., *Modelling and Using Context – Proceedings of the 2nd International and Interdisciplinary Conference, Context'99*, vol. 1688 of *Lecture Notes in Artificial Intelligence*, pages 145–158. Springer Verlag - Heidelberg.
- Ghidini, C. and F. Giunchiglia. 2001. Local models semantics, or contextual reasoning = locality + compatibility. *Artificial Intelligence* 127(2):221–259.
- Ghidini, C. and L. Serafini. 1998a. Information Integration for Electronic Commerce. In C. Sierra and P. Noriega, eds., *Proceedings of the Workshop on Agent Mediated Electronic Trading (AMET'98)*, vol. 1571 of *LNAI*, pages 189–206. Minneapolis, USA: Springer Verlag.
- Ghidini, C. and L. Serafini. 1998b. Model Theoretic Semantics for Information Integration. In F. Giunchiglia, ed., *Proceedings AIMS'98, 8th International Conference on Artificial Intelligence, Methodology, Systems, and Applications*, vol. 1480 of *LNAI*, pages 267–280. Sozopol, Bulgaria: Springer Verlag.
- Giunchiglia, E. and F. Giunchiglia. 1996. Ideal and Real Belief about Belief. In *Practical Reasoning, International Conference on Formal and Applied Practical Reasoning, FAPR'96*, no. 1085 in *Lecture Notes in Artificial Intelligence*, pages 261–275. Springer Verlag.
- Giunchiglia, F. 1993. Contextual reasoning. *Epistemologia, special issue on I Linguaggi e le Macchine* XVI:345–364. Short version in *Proceedings IJCAI'93 Workshop on Using Knowledge in its Context*, Chambéry, France, 1993, pp. 39–49.
- Giunchiglia, F. and L. Serafini. 1994. Multilanguage hierarchical logics, or: how we can do without modal logics. *Artificial Intelligence* 65(1):29–70.
- Giunchiglia, F., L. Serafini, E. Giunchiglia, and M. Frixione. 1993. Non-Omniscient Belief as Context-Based Reasoning. In *Proc. of the 13th International Joint Conference on Artificial Intelligence*, pages 548–554. Chambéry, France.
- Guha, R.V. 1991. *Contexts: a Formalization and some Applications*. Ph.D. thesis, Department of Computer Science, Stanford University. Also report No. STAN-CS-91-1399-Thesis.
- Lenat, D. B. 1995. Cyc: A large-scale investment in knowledge infrastructure. *Communications of the ACM* 38(11):33–38.
- McCarthy, J. 1987. Generality in Artificial Intelligence. *Communications of ACM* 30(12):1030–1035. Also in V. Lifschitz (ed.), *Formalizing common sense: papers by John McCarthy*, Ablex Publ., 1990, pp. 226–236.

REFERENCES / 23

- McCarthy, J. 1993. Notes on Formalizing Context. In *Proc. of the 13th International Joint Conference on Artificial Intelligence*, pages 555–560. Chambery, France.
- Mylopoulos, J. and R. Motschnig-Pitrik. 1995. Partitioning Information Bases with Contexts. In *Third International Conference on Cooperative Information Systems*. Vienna.
- Noriega, P. and C. Sierra. 1996. Towards Layered Dialogical Agents. In J. P. Müller, M. J. Wooldridge, and N. R. Jennings, eds., *Proceedings of the Third International Workshop on Agent Theories, Architectures, and Languages (ATAL-96), ECAI-96*, vol. 1193 of *LNAI*, pages 173–189. Budapest, Hungary.
- Parsons, S., C. Sierra, and N. R. Jennings. 1998. Agents that reason and negotiate by arguing. *Journal of Logic and Computation* 8(3):261–292.
- Theodorakis, M., A. Analyti, P. Constantopoulos, and N. Spyrtatos. 1998. 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: IEEE Computer Society.

November 8, 2002

2

On the dimensions of context dependence

MASSIMO BENERECETTI, PAOLO BOUQUET, CHIARA GHIDINI

In this chapter we propose to re-read the past work on formalizing context as the search for a logic of the relationships between partial, approximate, and perspectival theories of the world. The idea is the following. We start from a very abstract analysis of a context dependent representation into three basic elements. We briefly show that all the mechanisms of contextual reasoning that have been studied in the past fall into three abstract forms: *expand/contract*, *push/pop*, and *shifting*. Moreover we argue that each of the three forms of reasoning actually captures an operation on a different dimension of variation of a context dependent representation, *partiality*, *approximation*, and *perspective*. We show how these ideas are formalized in the framework of MultiContext Systems, and briefly illustrate some applications.

2.1 Introduction

In the last twenty years, the notion of context has become more and more central in theories of knowledge representation in Artificial Intelligence (AI). The interest in context is not limited to AI, though. It is discussed and used in various disciplines that are concerned with a theory of representation, such as philosophy, cognitive psychology, pragmatics, linguistics. Despite this large amount of work, we must confess that we are very far from a general and unifying theory of context. Even if we restrict the focus to theories of representation and

language, it is very difficult to see the relationship between different works on contextual reasoning. As an example, there are good pieces of work on utterance contexts, belief (and other intensional) contexts, problem solving contexts, and so on, but it is not clear whether they address different aspects of the same problem, or different problems with the same name.

In this chapter we propose to re-read the past work on context as the search for a logic of the relationships between partial, approximate, and perspectival theories of the world. The idea is the following. We start from an very abstract analysis of a context dependent representation into three basic elements: a collection of parameters (the contextual dependencies), a value for each parameter, and a collection of linguistic expressions (the explicit representation). Then, we briefly show that all the mechanisms of contextual reasoning that have been studied in the past fall into three abstract forms, *expand/contract*, *push/pop*, and *shifting*, each corresponding to an operation on one of the basic elements of the representation. Then, we argue that each of the three forms of reasoning actually captures an operation on a different dimension of variation of a context dependent representation, *partiality*, *approximation*, and *perspective*. This leads us to the conclusion that, at a suitable level of abstraction, *a logic of contextual reasoning is precisely a logic of the relationships between partial, approximate, and perspectival theories of the world*. We show how these ideas are formalized in the framework of MultiContext systems, and briefly illustrate some applications.

2.2 Contexts as boxes

In general, a representation is called context dependent when its content cannot be established by simply composing the content of its parts. In addition, one has to consider extra information that is left implicit in the representation itself. In Giunchiglia and Bouquet (1997), this notion of a context dependent representation is illustrated by introducing the so-called metaphor of the box (figure 1). A context dependent representation has three basic elements: a collection of parameters P_1, \dots, P_n, \dots , a value V_i for each parameter P_i , and a collection of linguistic expressions that provide an explicit representation of a state of affairs or a domain. The intuition is that the content of what is inside the box depends (at least partially, and in a sense to be defined) upon the values of the parameters associated with box. For example, in a context in which the speaker is John (i.e. the value of the parameter ‘speaker’ is set to John), the content (the intension, if you prefer) of the pronoun ‘I’ will be John, but this is not the case in a context in

which the speaker is Mary.

$$P1=V1 \dots Pn=Vn$$

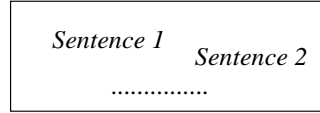


FIGURE 1 Contexts as boxes.

Starting from the metaphor of the box, it is quite easy to see that a theory of contextual reasoning is faced with a number of philosophical problems. A partial list includes: What features of context should be included among the parameters? Is it possible to specify *all* the relevant parameters, or the collection is always incomplete? How is the representation affected when the collection of parameters or their values changes? Can we get rid of parameters and get a context independent representation of the contents of a box? What is the relationship between the parameters of different boxes? How does this relationship affect the relationship between the contents of different boxes?

Since the goal of this chapter is not to provide a general foundation for a theory of context, we will not propose an answer to the issues above. Indeed, the analysis of the patterns of contextual reasoning is meant to hold no matter what solutions one adopts to these fundamental issues.

2.3 Forms of contextual reasoning

Mechanisms for contextual reasoning have been studied in different disciplines, though with different goals. A very partial list includes: *reflection* and *metareasoning* Weyhrauch (1980), Giunchiglia and Serafini (1994), *entering and exiting context, lifting, transcending context* Guha (1991), McCarthy (1993), Buvac and Mason (1993), *local reasoning, switch context* Giunchiglia (1993), Bouquet and Giunchiglia (1995), *parochial reasoning* and *context climbing* Dinsmore (1991), *changing viewpoint* Attardi and Simi (1995), *focused reasoning* Laird et al. (1987)). As a matter of fact, it is very difficult to see the relationship between these different works. We try to put some order in this situation by addressing the problem of identifying the general patterns of contextual reasoning, namely the general mechanisms that people use to reason with information (i) whose representation depend on a collection of contextual parameters, and (ii) which is scattered across a multiplicity of different contexts.

Our proposal is that all the forms of contextual reasoning that are discussed in the literature fall into three basic patterns, according to the element of the box that they affect: the representation, the collection of parameters, and the parameters' values.

Expand/Contract. A first general form of contextual reasoning (depicted in Figure 2) is based on the intuition that the explicit representation associated with a specific context does not contain all the facts potentially available to a reasoner, but only a subset of them. As a consequence, depending on the circumstances, the subset which is explicitly taken into account can be expanded (typically because some new input from the external environment makes it necessary to consider a larger collection of facts), or contracted (typically because the reasoner realizes that some facts are not relevant on a given occasion). An example of expansion is the Glasgow-London-Moscow (GLM) example McCarthy (1991), Bouquet and Giunchiglia (1995): when reasoning about traveling from Glasgow to Moscow via London, we normally do not include in the problem solving context the precondition that one must be dressed to get on a plane; however, if one's clothes are stolen at London airport, being clothed becomes a relevant precondition for the success of the travel plan, and therefore the original problem solving context must be expanded with facts about social conventions and buying clothes.

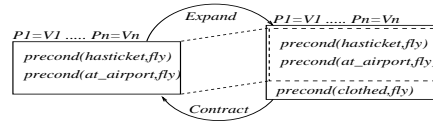


FIGURE 2 Expand/Contract.

In general, expansion and contraction are used to adjust a particular representation to a problem or to a given goal. The way problem solving contexts are built in CYC (using the strategy of lift-and-solve Guha (1991)), the mechanism of building appropriate mental spaces Fauconnier (1985) or partitioned representations Dinsmore (1991), and the process of selecting the relevant facts to interpret an utterance Sperber and Wilson (1986) are typical examples of this pattern of contextual reasoning.

Push/Pop. The content of a context dependent representation is partly encoded in the parameters outside the box, and partly in the sentences inside the box. Some authors propose reasoning mechanisms

for altering the balance between what is explicitly encoded inside the box and what is left implicit (i.e. encoded in the parameters). Intuitively, the idea is that we can move information from the collection of parameters outside the box to the representation inside the box, and vice versa. We call these two mechanisms *push* and *pop* to suggest a partial analogy with the operations of adding (pushing) and extracting (popping) elements from a stack. In one direction, *push* adds a contextual parameter to the collection outside the box and produces a flow of information from the inside to the outside of the box, that is part of what was explicitly encoded in the representation is encoded in some parameter. In the opposite direction, *pop* removes a contextual parameter from the collection outside the box and produces a flow of information from the outside to the inside, that is the information that was encoded in a parameter is now explicitly represented inside the box.

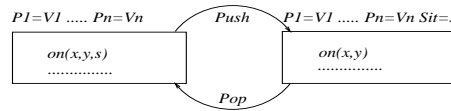


FIGURE 3 Push/Pop.

Consider, for instance, the well known AboveTheory scenario, introduced in McCarthy (1993). The fact that block x is on block y in a situation s is represented as $on(x, y, s)$ in a context c with no parameter for situations. This is because in some cases we want to leave implicit the dependence on the situation s (typically, when we don't want to take situations into account in reasoning). This means that the situation can be encoded as a parameter, and the representation can be simplified to $on(x, y)$. Push is the reasoning mechanism which allows us to move from $on(x, y, s)$ to $on(x, y)$ (left-to-right arrow in figure 3), whereas pop is the reasoning mechanism which allows us to move back to $on(x, y, s)$ (right-to-left arrow in figure 3). Hence, push and pop capture the interplay between the collection of parameters outside the box and the representation inside the box.

It is worth noting that the mechanism of entering and exiting context proposed by McCarthy and others can be viewed as an instance of push and pop. Suppose we start with a sentence such as $c_0c : p$, whose intuitive meaning is that in context c_0 it is true that in context c the proposition p is true. The context sequence c_0c can be viewed as the reification of a collection of parameters. Exiting c pops the context sequence, and the result is the formula $c_0 : ist(c, p)$, where the dependence

on c is made explicit in the representation $ist(c, p)$ ($ist(c, p)$ is the main formula of McCarthy's formalism, asserting that a p is true in context c); conversely, entering c pushes the context sequence and results in the formula $c_0c : p$, making the dependence on c implicit in the context sequence. Other examples of push/pop are: *reflection up* to pop the collection of parameters and *reflection down* to push it in Giunchiglia and Serafini (1994); the rule of *context climbing* to pop the collection of parameters, and the rule of *space initialization* to push it in Dinsmore (1991).

Shifting. Shifting changes the value of one or more contextual parameter, without changing the collection of parameters. The name 'shifting' is inspired to the concept of shifting in Lewis (1980). The intuition is that changing the value of the parameters shifts the interpretation of what is represented inside the box.

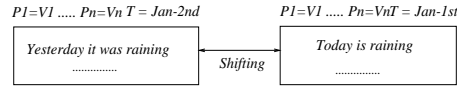


FIGURE 4 Shifting.

The simplest illustration of shifting is again indexical expressions. The fact that on January 1st it is raining is represented as 'Today is raining' in a context in which time is set to January 1st, but it is represented as 'Yesterday it was raining' if the value of time changes to January 2nd. As it is shown in figure 4, *shifting* is the reasoning mechanism which allows us to move from one representation to the other by changing the value of the parameter time, provided we know the relationship between the two parameter's values. Another very common example of shifting is when the viewpoint changes, e.g. when two people look at the same room from opposite sides (what is right for the first will be left for the other). A third case is categorization. For the supporters of team A, the members and the supporters of team B are opponents, and vice versa for the supporters of team B. And the examples can be multiplied.

In the literature, we can find different instances of shifting. Kaplan's notion of *character* is the semantical counterpart of this reasoning mechanism with indexical languages; Guha and McCarthy formalize a form of shifting using the notion of *lifting* Guha (1991); Dinsmore introduces the notion of *secondary context*.

2.4 Dimensions of context dependence

Our next step is to show that the three forms of contextual reasoning actually operate each on a fundamental dimensions of a context dependent representation: *partiality*, *approximation*, and *perspective*. We start with a more precise characterization of partiality, approximation, and perspective.

Partiality. We say that *a representation is partial when it describes only a subset of a more comprehensive state of affairs*. We observe that the notion of partiality can be analyzed from two different perspectives: metaphysically, a theory is partial if it does not cover the entire universe; however, cognitively, a representation is partial if it does not cover the totality of what an agent can talk about. For our present purposes, either perspective is acceptable, even though our general attitude is in favor of the cognitive view.

Perhaps the more intuitive example of partial theories are domain specific theories. For instance, a theory about the theory about the Italian cuisine is partial because it does not provide information about Indian or French cuisine, about soccer, about quantum mechanics. A different usage of partial theories is in problem solving. Given a problem, people seem to be capable of circumscribing what knowledge is relevant to solve it, and disregard the rest. In this case, assumptions on what is relevant act as contextual parameters. Partial theories are also used in theories of linguistic communication. When a speaker says something to a hearer, it is assumed that the latter interprets what the speaker said in some context. According to Sperber and Wilson (1986), '[a] context is a psychological construct, a subset of the hearer's assumptions about the world'. Such a context includes the set of facts that the hearer takes to be relevant in order to assign the correct interpretation to what the speaker said. In this sense, it is a partial theory.

Partiality is a relative notion. Intuitively, there is a partial order between partial representations. Therefore a representation can be more or less partial of another one. Two partial representations may also overlap. We do not further discuss these aspect here. We only need to make clear the idea that partiality is a dimension along which a representation may vary.

Approximation. We say that *a representation is approximate when it abstracts away some aspects of a given state of affairs*. A representation of the blocks world in terms of the binary predicates $on(x, y)$ e $above(x, y)$ is approximate, because the time (situation) is abstracted away.

As for partiality, approximation is a relative notion: a representation

is approximate because it abstracts away details that another representation takes into account. The representation $on(x, y)$ and $above(x, y)$ is more approximate than the representation $on(x, y, s)$ and $above(x, y, s)$ because the first abstracts away the dependence on the situation. Of course, an open point is whether there is such a thing as a non approximate representation of a state of affairs. This would be a sort of least approximate representation, namely a representation which is less approximate than anyone else. We avoid committing to one position or the other; here we are interested in the reasoning mechanisms that allow us to switch from a more to a less approximate representation (and vice versa), and not in the epistemological status of representations.

Perspective. A third dimension along which a representation may vary is perspective. We say that *a representation is perspectival when it encodes a spatio-temporal, logical, or cognitive point of view on a state of affairs.*

The paradigmatic case of spatio-temporal perspective is a given by indexical languages. A sentences such as ‘It’s raining (here)(now)’ is a perspectival representation because it encodes a spatial perspective (i.e. the location at which the sentences are used, the speaker’s current ‘here’) and a temporal perspective (i.e. the time at which the sentences are used, the speaker’s current ‘now’). The philosophical tradition shows us that some sentences (e.g. ‘Ice floats on water’) encode a logical perspective as well, because they implicitly refer to ‘this’ world, namely the world in which the ‘here’ and ‘now’ of the speaker belong (the same sentence, if uttered in a world different from our world, might well be false). Thus Kaplan includes a world among the features that define a context, and uses this world to interpret the propositional operator ‘actually’.

Indexicals are not the only type of expressions that encode a physical perspective. Suppose, for example, that two agents look at the same object (for example the magic box of figure 5). Because of their different viewpoints, the representation of what they see is completely different, and the same ball can be described as being on the right by **Side** and as being on the left by **Front**.

A subtler form of perspective is what we call cognitive perspective. It has to do with the fact that many representation encode a point of view which includes a collection of beliefs, intentions, goals, and so on. Cognitive perspective is very important in the analysis of what is generally called an *intensional context*, such as a belief context. John and Mary may have dramatically different beliefs about Scottish climate, even if they represent the same universe of discourse (or portion of the

world) at the same level of approximation. We don't see any other way of making sense of this difference than that of accepting the existence of a cognitive perspective, which is part of what determines the context of a representation.

At this point, we are ready to justify our claim that the three forms of contextual reasoning are precisely mechanisms that operate on the dimensions of partiality, approximations, and perspective:

- Expand/Contract is the reasoning mechanism that allows us to vary the degree of partiality by varying the amount of knowledge which is used in the representations of the world.
- Push/Pop is the reasoning mechanism that allows us to vary the degree of approximation by regulating the interplay between the collection of parameters outside and the explicit representation inside a box.
- Shifting is the reasoning mechanism that allows us to change the perspective by taking into account the 'translation' of a representation into another when the value of some contextual parameter is changed.

As a consequence of our claim, a logic of contextual reasoning must formalize the reasoning mechanisms of expand/contract push/pop, and shifting and use them to represent the relationship between partial, approximate, and perspectival representations. Our final step is to show that MultiContext systems satisfy this requirement, and to validate this assertion by analyzing in detail some applications of MultiContext systems.

2.5 A logic of contextual reasoning: MultiContext Systems

In the past, various logics have been proposed which formalize one aspect or the other of such a logic of contextual reasoning. For example, Kaplan's logic of demonstratives is a logic which allows only for a multiplicity of perspectival representations (partiality and approximation are left unchanged from one context to the other) and, consequently, provides only mechanisms for shifting (in the form of the semantic notion of character). Buvač and Mason's propositional logic of context allows for a multiplicity of partial, approximate, and perspectival representations, and provides the machinery for expand/contract, push/pop, and shifting; however, it formalizes a quite weak form of partiality (via the use of partial functions for interpreting a global language) and only a special form of push/pop (i.e. making explicit or implicit the context itself).

MultiContext systems (MCS) Giunchiglia and Serafini (1994), and their *Local Model Semantics* (LMS) Ghidini and Giunchiglia (2001), provide a logic for contextual reasoning based on the principles of *locality* and *compatibility*. These principles are stated and discussed in Chapter 1, where LMS is introduced and explained in detail. For the sake of our presentation discussion, we rewrite these principles as follows:

1. each context c_i is associated with a different formal language L_i , used to describe what is true in that context. The semantics of L_i is *local* to the context itself. Therefore each context has its own set of local models M_i , and local satisfiability relation \models_i ;
2. the relations between different contexts are modeled by means of *compatibility* relations between (sets of) local models of the different contexts.

We believe that principle of locality and principle of compatibility provide LMS and MCS with the capability of being a suitable logic of the relation between partial, approximate, and perspectival representations. For lack of space, we focus the discussion of our claim on LMS. A similar analysis applies to the axiomatic system MCS.

By associating distinct languages and local semantics to different contexts, LMS allows for different partial, approximate, and perspectival representations. The most intuitive case is partial representations. Simple examples are: the language might contain only a subset of a more comprehensive set of symbols, the class of models might satisfy only a subset of a more comprehensive set of axioms, or rules of well-formedness. Second, approximate representations. A simple case is the AboveTheory example: a context might contain the binary predicate $on(x, y)$ or the ternary predicate $on(x, y, s)$ depending on the fact that the it abstracts away or represents the dependence on the situation. Third, perspectival representations. An example is the fact that the truth value of a formula in a context might depend on the perspective from which the world is represented. The truth value of the formula might change in different contexts, depending on the corresponding shift of perspective.

By modeling compatibility relations between different contexts as relations among the (sets of) local models of the different contexts LMS allows one to represent the relations existing between a multiplicity of partial, approximate, and perspectival representations of the world. For instance, if the relation contains a pair $\langle models(c_1), models(c_2) \rangle$ composed by a set $models(c_1)$ of models of context c_1 and a set $models(c_2)$ of models of context c_2 , and all the models in $models(c_1)$ are obtained

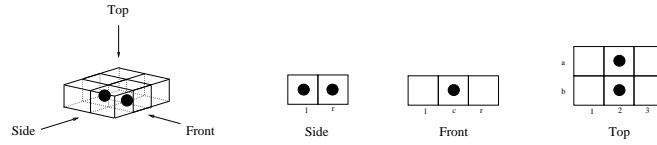


FIGURE 5 The magic box and its partial views.

as the expansion of a model in $models(c_2)$, then it is easy to observe that c_2 describes a portion of the world which is a subset of the portion described by c_1 . By studying and classifying the different relations existing among the (sets of) local models of the different contexts we might, in principle, try to classify the many different relations existing among different partial, approximate, and perspectival representations. Unfortunately, even if we restrict ourselves to model each context c_i by mean of a first order language and the classical semantics, we must admit that we are still far from having a (even partial) classification of these many different relations. Although some of them are very easy to identify, as the relation of expansion mentioned above, relations between partial, approximate and perspectival representations may be, in general, much trickier. Nonetheless, by analyzing existing applications of LMS and MCS we are able to show that LMS and MCS have been used to represent context-based representation and reasoning in terms of the relations among partial, approximate, and perspectival representations. In the rest of the section briefly show the result of our analysis. This provides a first evidence of the fact that LMS is a logic of the relations between partial, approximate and perspectival representations. This provides also a first motivation for a future work on studying and classifying the many different relations existing among different partial, approximate, and perspectival representations.

Viewpoints. A paradigmatic example of reasoning with viewpoint is the Magic Box (MB) example, developed in Benerecetti et al. (2000).

There are three observers, **Top**, **Side**, and **Front**, each having a partial view of a box as shown in the top part of Figure 5. **Top** sees the box from the top, and **Side** and **Front** see the box from two different sides. The box consists of six sectors, each sector possibly containing a ball. The box is “magic” and **Side** and **Front** cannot distinguish the depth inside it. The bottom part of Figure 5 shows the views of the three agents corresponding to the scenario depicted in the top part. **Top**, **Side**, and **Front** decide to test their new computer program ϵ by submitting the following puzzle to it. **Side** and **Front** tell ϵ their partial views. Then they ask ϵ to guess **Top**’s view of the box.

Benerecetti et al. (2000) describes a formalization of the reasoning process of ϵ in solving the puzzle, by mean of the four contexts depicted in figure 6. Contexts **Side** and **Front** contain the program's

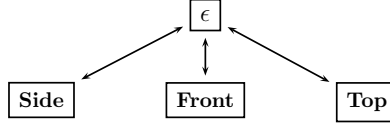


FIGURE 6 The four contexts in the MB example.

representation of **Side**'s and **Front**'s knowledge; context **Top** contains the program's representation of **Top**'s knowledge, and is the context in which it will try to build the solution; context ϵ contains the knowledge that the computer program has about the game, namely what the relations among the other contexts are.

According to our classification of dimensions of a context dependent representation, the representations of the different contexts **Side** and **Front**, **Top**, and ϵ may vary along three dimensions: partiality, approximation, and perspective. Focusing on partiality, the different contexts are related to different specific domains. For instance, **Side** can only talk about the (non) presence of a ball in the left or right sector it sees, **Front** can talk about the (non) presence of a ball in the left, or the central or right sector it sees, **Top** can talk about the presence of a ball in each one of the six sectors, while ϵ needs only to talk about how the pieces of knowledge contained in each one of the contexts above are related to each other. Focusing on approximation, we notice that the description of (a portion of) the world in **Side**, **Front**, and **Top** is given in terms of balls and sectors of the box, whereas the description in context ϵ concerns how to relate the information coming from the different observers. In order to do this, context ϵ needs to make explicit some information that was implicit in the observers' contexts. In particular, it needs to make explicit what information comes from what observer. This is an example of push/pop and is related to the different levels of approximation of the different contexts. In this case we say that the representation in **Side**, **Front**, and **Top** is more approximate than the one in ϵ , because the first ones abstract away what information comes from what observer. Focusing on perspective, each of the observer's contexts expresses knowledge about the box which depends on the observer's physical perspective. For example, the fact that **Side** sees a ball in the left sector (from his point of view) is different from **Front** seeing a ball in the left sector (from his point of view). Since their

perspectives are different, the same description (e.g., ‘A ball is in the left sector’) may, thus, have a different meaning in different contexts.

Formally, the specific domains of **Side**, **Front**, and **Top** are described by three different propositional languages L_{Side} , L_{Front} and L_{Top} built up from the sets $AP_{\text{Side}} = \{l, r\}$, $AP_{\text{Front}} = \{l, c, r\}$, and $AP_{\text{Top}} = \{a1, a2, a3, b1, b2, b3\}$ of propositional constants (where l means that the observer sees a ball in the *left* sector, c means that the observer sees a ball in the *central* sector, and so on) To account for the specific domain of ϵ , and its shift in the approximation level described above, the language L_ϵ contains a set $\{Side, Front, Top\}$ of constant symbols for each one of the contexts above, a set of constant symbol “ ϕ ” for each formula ϕ that can be expressed in the languages L_{Side} or L_{Front} or L_{Top} , and a binary predicate $ist(c, “\phi”)$, whose intuitive meaning is that formula $\phi \in L_c$ is true in context c (see McCarthy (1993)). In order to solve the puzzle ϵ needs to relate information contained in different contexts associated with different levels of approximations. In particular Benerecetti et al. (2000) needs to formalize the relation denoted as arrows connecting contexts in figure 6. This is done by imposing a compatibility relation between the models of each observers’ context c and models of the context ϵ . To state the correspondence between a formula ϕ in each observers’ context c and the formula $ist(c, “\phi”)$ (denoting the same fact at a different degree of approximation) in context ϵ , if a formula of the form $ist(c, “\phi”)$ is a theorem in ϵ , then the formula ϕ must be a theorem in c , and vice-versa. The different perspectival representations are formalized in Benerecetti et al. (2000) by the different (initial) axioms satisfied in each observer’s context and the relations between them are explicitly stated as axioms in context ϵ .

Belief contexts. LMS and MCS have been applied to formalize different aspects of intentional contexts, and in particular belief contexts (see e.g., Giunchiglia et al. (1993), Cimatti and Serafini (1995)). An example is a puzzle described in Benerecetti et al. (1998), where LMS and MCS are used to solve the problem of the opaque and transparent reading of belief reports.

A computer program ϵ knows that Mr. A believes that the president of the local football team is Mr. M and that Mr. B believes that the president is Mr. C . The computer program knows also that Mr. B knows that A believes that the president of the local football team is Mr. M . Actually, Mr. B is right, and the computer program knows that. Now, B tells ϵ : “Mr. A believes that the president of the local football team is a corruptor”. How will ϵ interpret the sentence?

The program is a little puzzled, since the question has two possible answers: (1) Mr. *A*'s belief is referred to Mr. *M* (since Mr. *A* is the subject of the belief). This is an instance of opaque reading. (2) Mr. *A*'s belief is referred to Mr. *C* (since it is Mr. *B* who is speaking). This is an instance of transparent reading.

We are not interested here in the solution of the puzzle (the interested reader may refer to Benerecetti et al. (1998)), but in analyzing the representations of the different contexts involved in the formalization.

The formalization is based on the notion of *belief context*. A belief context is a representation of a collection of beliefs that a reasoner (in this example, the program) ascribes to an agent (including itself) from a given perspective. Possible perspectives are: the beliefs that the program ascribes to itself (e.g., that Mr. *B* believes that Mr. *A* believes that the president of the local football team is a corruptor); the beliefs that the program ascribes to Mr. *B* (e.g., that Mr. *A* believes that the president of the local football team is a corruptor); the beliefs that the program ascribes to Mr. *B* about Mr. *A* (e.g., that the president of the local football team is a corruptor). The belief contexts that the program can build in this example can be organized in a structure like that presented in figure 7.

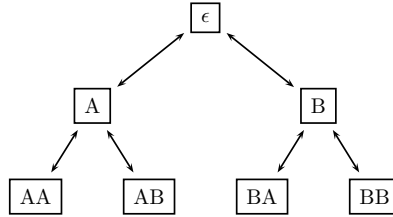


FIGURE 7 The structure of belief contexts

ϵ represents the context containing the beliefs that the program ascribes to itself, **A** is the context containing the beliefs that the program ascribes to Mr. *A*, **B** is the context containing the beliefs that the program ascribes to Mr. *B*, **BA** is the context containing the beliefs that the program ascribes to Mr. *A* from Mr. *B*'s perspective, and so on.

The formalization of the different contexts in figure 7 may vary along the three dimensions of contextual dependence. Focusing on partiality, the different contexts are related to different sets of beliefs. For instance, **A** is the context containing the beliefs that the program ascribes to Mr. *A*, whereas **B** is the context containing the beliefs that the program ascribes to Mr. *B*. Focusing on approximation, we notice that each

context in the hierarchy must be able to talk about beliefs contained in each one of the contexts above. In order to do this it needs to make explicit some information that was implicit in the observers' contexts. In particular, it needs to make explicit what beliefs are contained in what context. The relations involving different contexts associated with different degrees of approximations are the one denoted as arrows in figure 6 and are similar to the ones described in the MB example. Focusing on perspective, each of the belief contexts expresses knowledge about the world which depends on the cognitive perspective of the agents, from the point of view of the computer program. For instance, Mr. *B* will refer to Mr. *C* as “the president of the local football team”, whereas Mr. *A* will refer to Mr. *C* as Mr. *M*.

Integration of different information sources. LMS and MCS have been applied to formalize the integration of information coming from different information sources. Ghidini and Serafini (1998a,b) contain the formal definitions and motivating examples. Let us focus on a simple example.

A mediator *m* of an electronic market place collects information about fruit prices from 1, 2, and 3 and integrates it in a unique homogeneous database. Customers that need information about fruit prices may therefore submit a single query to the mediator instead of contacting the sellers.

The formalization of the exchange of information in this example based on the four contexts and the information flows depicted in figure 8. Circles represent contexts associated to the different databases and arrows represent information flow between contexts (databases).

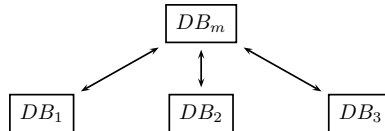


FIGURE 8 Contexts in the mediator example.

The representations of the different contexts in figure 8 may have different degrees of partiality, as each database is associated to a specific domain. For instance, the sellers might provide different subsets of fruits and therefore the domains of their databases are different. Focusing on approximation, the domain of fruits can be represented at different level of details by different sellers. E.g., database 1 may contain prices for red apples and yellow apples, while database 2 and 3 abstracts away the dependence on the color and do not make this distinction. Focusing on

perspective, prices of the different sellers might be not homogeneous, depending on their particular viewpoint. E.g., prices of database 1 don't include taxes, while prices of database 2, 3 and the mediator do.

Formally, the specific domains of the different databases are described by using different first order languages. Each database is associated with a different interpretation domain. The compatibility relation between the different levels of approximations in the fruit domains is formalized by using *domain relations*, i.e. relations between the interpretation domains of the different databases. A domain relation may, for instance, relate a “more abstract” object (e.g. apple) in the domain of a database to a set of “less abstract” objects (e.g. red-apple, green-apple) in the domain of another database. Compatibility relations between the different perspectival views contained in the databases are formalized by using *view constraints*, i.e. relations between formulae contained in different languages (databases). For instance every time the models of database 1 satisfy the formula $hasprice(x, y)$ (meaning that item x has price y , then the models of the mediator database must satisfy the formula $\exists y' hasprice(x, y') \wedge y' = y + (0.07 * y)$ (meaning that the same item x has price y' which is obtained adding the amount of taxes to y).

2.6 Conclusions

In the chapter, we have not presented a new theory about partiality, approximation, or perspective. Instead, we have shown that the work on contextual reasoning in AI (and not only in AI) can be re-read as an attempt of providing a logic of the mechanisms that govern the relationship between partial, approximate, and perspectival representations of the world.

In this sense, the work described here is only a preliminary step. Indeed, it opens a whole field of research, both philosophical and logical. Our next step will be a formal study of a logic of partiality, approximation, and perspective in the framework of LMS and MCS. In particular, we are interested in finding the compatibility relations (and the relative bridge rules) involved in the corresponding reasoning mechanisms. This, we hope, will be part of a new approach to a theory of knowledge representation, in which context will play a crucial role.

References

- Attardi, G. and M. Simi. 1995. A formalisation of viewpoints. *Fundamenta Informaticae* 23(2-4):149-174.
- Benerecetti, M., P. Bouquet, and C. Ghidini. 1998. Formalizing Belief Reports – The Approach and a Case Study. In F. Giunchiglia, ed., *Proceed-*

REFERENCES / 41

- ings AIMS'98, 8th International Conference on Artificial Intelligence, Methodology, Systems, and Applications, vol. 1480 of *LNAI*, pages 62–75. Sozopol, Bulgaria: Springer Verlag.
- Benerecetti, M., P. Bouquet, and C. Ghidini. 2000. Contextual Reasoning Distilled. *Journal of Theoretical and Experimental Artificial Intelligence* 12(3):279–305.
- Bouquet, P. and F. Giunchiglia. 1995. Reasoning about Theory Adequacy. A New Solution to the Qualification Problem. *Fundamenta Informaticae* 23(2–4):247–262. Also IRST-Technical Report 9406-13, IRST, Trento, Italy.
- Buvac, Sasa. and Ian A. Mason. 1993. Propositional logic of context. *Proc. of the 11th National Conference on Artificial Intelligence*.
- Cimatti, A. and L. Serafini. 1995. Multi-Agent Reasoning with Belief Contexts: the Approach and a Case Study. In M. Wooldridge and N. R. Jennings, eds., *Intelligent Agents: Proceedings of 1994 Workshop on Agent Theories, Architectures, and Languages*, no. 890 in Lecture Notes in Computer Science, pages 71–85. Springer Verlag. Also IRST-Technical Report 9312-01, IRST, Trento, Italy.
- Dinsmore, J. 1991. *Partitioned Representations*. Kluwer Academic Publishers.
- Fauconnier, G. 1985. *Mental Spaces: aspects of meaning construction in natural language*. MIT Press.
- Ghidini, C. and F. Giunchiglia. 2001. Local models semantics, or contextual reasoning = locality + compatibility. *Artificial Intelligence* 127(2):221–259.
- Ghidini, C. and L. Serafini. 1998a. Distributed First Order Logics. In D. Gabbay and M. de Rijke, eds., *Frontiers Of Combining Systems 2*, Studies in Logic and Computation, pages 121–140. Research Studies Press. ISBN 0-86380-252-4.
- Ghidini, C. and L. Serafini. 1998b. Information Integration for Electronic Commerce. In C. Sierra and P. Noriega, eds., *Agent Mediated Electronic Commerce. First International Workshop on Agent Mediated Electronic Trading, AMET-98*, vol. 1571 of *LNAI*, pages 189–206. Minneapolis, USA: Springer Verlag. Also accepted for presentation at UKMAS'98. Manchester, U.K. December 14–15, 1998.
- Giunchiglia, F. 1993. Contextual reasoning. *Epistemologia, special issue on I Linguaggi e le Macchine* XVI:345–364. Short version in Proceedings IJCAI'93 Workshop on Using Knowledge in its Context, Chambéry, France, 1993, pp. 39–49. Also IRST-Technical Report 9211-20, IRST, Trento, Italy.
- Giunchiglia, F. and P. Bouquet. 1997. Introduction to contextual reasoning. An Artificial Intelligence perspective. In B. Kokinov, ed., *Perspectives on Cognitive Science*, vol. 3, pages 138–159. Sofia: NBU Press. Lecture Notes of a course on “Contextual Reasoning” of the European Summer School on Cognitive Science, Sofia, 1996.

- Giunchiglia, F. and L. Serafini. 1994. Multilanguage hierarchical logics, or: how we can do without modal logics. *Artificial Intelligence* 65(1):29–70. Also IRST-Technical Report 9110-07, IRST, Trento, Italy.
- Giunchiglia, F., L. Serafini, E. Giunchiglia, and M. Frixione. 1993. Non-Omniscient Belief as Context-Based Reasoning. In *Proc. of the 13th International Joint Conference on Artificial Intelligence*, pages 548–554. Chambery, France. Also IRST-Technical Report 9206-03, IRST, Trento, Italy.
- Guha, R.V. 1991. Contexts: a Formalization and some Applications. Tech. Rep. ACT-CYC-423-91, MCC, Austin, Texas.
- Laird, J. E., A. Newell, and P.S. Rosenbloom. 1987. Soar: An architecture for general intelligence. *Artificial Intelligence* 33(3):1–4 64.
- Lewis, D. 1980. Index, Context, and Content. In S. Kranger and S. Ohman, eds., *Philosophy and Grammar*, pages 79–100. D. Reidel Publishing Company.
- McCarthy, J. 1991. Overcoming an Unexpected Obstacle. Unpublished.
- McCarthy, J. 1993. Notes on Formalizing Context. In *Proc. of the 13th International Joint Conference on Artificial Intelligence*, pages 555–560. Chambery, France.
- Sperber, Dan and Deirdre Wilson. 1986. *Relevance. Communication and Cognition*. Basil Blackwell.
- Weyhrauch, R.W. 1980. Prolegomena to a Theory of Mechanized Formal Reasoning. *Artificial Intelligence* 13(1):133–176.