Local Models Semantics, or Contextual Reasoning = Locality + Compatibility

Fausto Giunchiglia[1,2] Chiara Ghidini[2]
[1]IRST, Povo, 38100 Trento, Italy
[2]DISA - University of Trento, Via Inama 5, 38100 Trento, Italy
fausto@irst.itc.it, ghidini@cs.unitn.it

Abstract

In this paper we present a new semantics, called *Local Models Semantics*, and use it to provide a foundation to reasoning with contexts. This semantics captures and makes precise the two main intuitions underlying contextual reasoning: *(i)* reasoning is mainly *local* and uses only part of what is potentially available (e.g. what is known, the available inference procedures), this part is what we call *context* (of reasoning); however *(ii)* the context may change and there is *compatibility* among the reasoning performed in different contexts. We validate our semantics by formalizing two important forms of contextual reasoning: reasoning with viewpoints and reasoning about belief.

1 Introduction

Contexts are an important topic in many research areas, for example in the semantics of natural language (see, e.g., [Kaplan, 1978]), and in computational linguistics (see, e.g., [Fauconnier, 1985]). Lately, contexts have been independently proposed in [Giunchiglia, 1993] and [McCarthy, 1987] as an important means for formalizing (certain forms of) reasoning. In [Giunchiglia, 1993] contexts are seen as a tool for formalizing the locality of reasoning, while in [McCarthy, 1987] contexts are seen as a way for solving the problem of generality. Coherently with these two proposals, contexts have been used in various applications, e.g. in the integration of heterogeneous knowledge and data bases [Farquhar *et al.*, 1995; Mylopoulos and Motschnig-Pitrik, 1995], in the formalization of reasoning about beliefs [Giunchiglia, 1993; Giunchiglia and Serafini, 1994], and in the formalization of reasoning with viewpoints [Attardi and Simi, 1995].

Despite the pletora of different approaches, formalizations, and applications, two are the main intuitions underlying the use of context in reasoning (and also in the other research areas). We state these two intuitions 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): the context may change. However, there is *compatibility* among the reasoning performed in different contexts.

The goal of this paper is to describe and motivate a new semantics, called *Local Models Semantics*, which formalizes the two principles listed above, and that we propose as a foundation of contextual reasoning. In Section 2 we give two motivating examples. In Section 3 we present Local Models Semantics and discuss how it formalizes the principles of locality and compatibility, plus other intuitions. In Section 4 we formalize the two examples introduced in Section 2. Finally, in Section 5 we compare Local Models semantics with other frameworks proposed in the past. We end with some concluding remarks (Section 6).

2 Two motivating examples

The first example is about reasoning with viewpoints, the second is about reasoning about belief. These examples highlight how the semantics and methodology developed in this paper can be applied, suitably generalized, to the modelling of two important problems: the federation of heterogeneous data or knowledge bases (first example) and multiagent systems (second example).

2.1 Viewpoints

Consider the scenario of Figure 1. There are two observers, $Mr.1$ and $Mr.2$, each having a partial view of a box. The box consists of six sectors, each sector possibly containing a ball. There cannot be balls hidden from the view of an observer. The box is "magic" and the two observers cannot distinguish the depth inside it. We have two contexts, each context for what an observer sees and the consequences that it is able to draw from it.

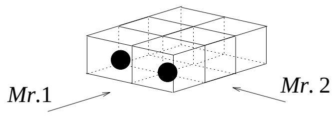

Figure 1: The magic box.

Locality: Both $Mr.1$ and $Mr.2$ have the notions of a ball being on the right or on the left. However these two notions are different and we may have a ball which is on the right for $Mr.1$ and not on the right for $Mr.2$. Furthermore $Mr.2$ has the notion of "a ball being in the center of the box" which is meaningless for $Mr.1$.

Compatibility: $Mr.1$ and $Mr.2$'s views are viewpoints of the same box and, therefore, there must be a correspondence between them. Figure 2 gives all the views and possible combinations of compatible views of $Mr.1$ and $Mr.2$. We can describe this situation by listing all the possible compatible pairs (exactly as they are represented in Figure 2), or we can describe it more synthetically, and possibly also less precisely. For instance we can have a description like: "if $Mr.1$ sees at least a ball then $Mr.2$ sees at least a ball".

Notice that the most straightforward formalization would be a direct axiomatization of the box as a two-dimensional grid. $Mr.1$ and $Mr.2$'s views could then be

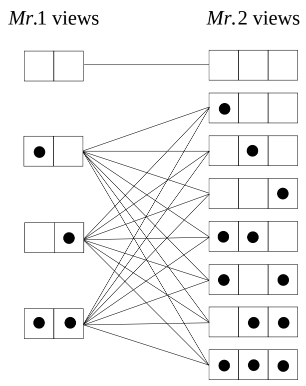

Figure 2: Compatible views of $Mr.1$ and $Mr.2$.

easily 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 we can only reconstruct a partial compatibility among them. An important application domain where this may or may not be the case is the development and integration of data or knowledge bases. Thus 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. The work in [Serafini and Ghidini, 1997] starts from this observation, further develops the semantics defined in this paper, and gives foundations to the various kinds of federation described in, e.g., [Farquhar *et al.*, 1995; Mylopoulos and Motschnig-Pitrik, 1995].

2.2 Reasoning about belief

Consider now the situation of two agents, a and b with a having beliefs about the beliefs of b (see Figure 3). We have again two contexts, one for what a believes and one for what b believes.

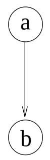

Figure 3: a has beliefs about the beliefs of b.

Locality: a and b have different beliefs about the world. Furthermore, a has the notion of "b believing something" which b does not have.

Compatibility: The level of compatibility depends on the situation. We may, for instance, have that a's beliefs about b's beliefs are correct, in which case if a believes that b believes ϕ then b believes ϕ. Another situation is when a's beliefs about b's beliefs are complete, in which case if b believes ϕ then a believes that b believes ϕ. A taxonomy of the possible forms of belief about belief is introduced in [Giunchiglia *et al.*, 1993].

The presentation of a multiagent system as a set of contexts, each context representing an agent, was first advocated in [Giunchiglia *et al.*, 1993]. In this approach the reasoning capabilities of an agent, its interactions with the other agents, and its beliefs about the beliefs of the other agents can be directly modelled as properties of the context representing the agent. A specification given in this way has properties of modularity and incrementality which are missing in other frameworks. The specifications developed in [Giunchiglia *et al.*, 1993] were axiomatic and set-theoretic. This paper highlights how this approach can be given a semantic foundation which maintains modularity and incrementality.

3 Local Models Semantics

For the sake of simplicity, let us consider the situation with two contexts (the generalization to any enumerable set of contexts does not present conceptual difficulties). Let us restrict ourselves to (classes of) first order models and let us suppose that L_1, L_2 are the two first order languages used to describe what is true in the two contexts. Let $\overline{M}_i$ be the class of all the models (interpretations) of L_i. Let $\mathbf{C}$,

$$\mathbf{C} \subseteq 2^{\overline{M}_1} \times 2^{\overline{M}_2},$$

be a *compatibility relation* (for $\{L_1, L_2\}$). Let us call $m \in \overline{M}_i$ a *local model* (of L_i) and $\mathbf{c} \in \mathbf{C}$ a *(compatibility) pair* (for $\{L_1, L_2\}$). Notationally, if $\mathbf{c}$ is a pair, we write $\mathbf{c}_1$ and $\mathbf{c}_2$ to mean its first and second element, respectively; and $\mathbf{c}_i$ to mean its i-th element, with i possibly equal to 1 or 2.

A *model* is a compatibility relation which contains at least a pair and does not contain the pair of two empty sets.

Definition 3.1 (Model) *A* model *(for $\{L_1, L_2\}$) is a compatibility relation $\mathbf{C}$ such that:*

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

In the following we write $\mathbf{C}$ to mean either a compatibility relation or a model. The context always makes clear what we mean. Figure 4 gives a graphical representation of the construction we perform. We start from L_1, L_2. Then, we associate each of them with a set $M_i \subseteq \overline{M}_i$ of local models. Usually $M_i \subset \overline{M}_i$. For instance, in the magic box scenario, if there is only one ball in the box, then all the local models (for both $Mr.1$ and $Mr.2$) which allow zero, two or three balls must be discarded. Finally, we pair local models inside compatibility pairs. The resulting compatibility relation is our model. For instance, we can construct a local model for each view in Figure 2 and then build a compatibility relation which contains the pairs constructed as follows: as first element take a singleton set containing a model formalizing a view of $Mr.1$, as second element take a singleton set containing a model formalizing a view compatible with the chosen view of $Mr.1$. Notice that linking local models inside a compatibility relation may force us to throw away some of them. For instance, considering the local models for $Mr.1$ which allow for only one ball forces us to throw away all the pairs and corresponding local models for $Mr.2$ which allow for zero balls.

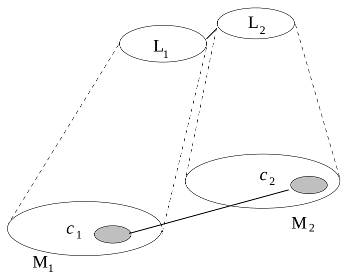

Figure 4: Model for $\{L_1, L_2\}$.

A *context* C_i, with $i = 1, 2$, (in a model $\mathbf{C}$) is (formally defined as) the set of local models $m \in \overline{M}_i$ allowed by $\mathbf{C}$, in symbols

$$C_i = \{m \ : \ m \in \mathbf{c}_i \text{ with } \mathbf{c} \in \mathbf{C}\}.$$

Notice that defining a context as a set of models (instead of, e.g., a single model) allows us to formalize it as a partial object, as explicitly required in, e.g., [Giunchiglia, 1993; McCarthy, 1987]. This is a key difference with possible worlds, which are complete objects (in the sense that a formula is either true or false in a world). The advantages of having contexts as partial objects are easily seen. Consider for instance the example of Section 2.2. In this case we may impose that a is uncommitted to whether b believes ϕ simply by having ϕ true in some local models of the context for a and false in others. A similar observation applies to compatibility pairs. For instance, in the example of Section 2.1, we can impose a compatibility constraint which says that if $Mr.1$ sees at least a ball then $Mr.2$ sees at least a ball. In this case we have in the same compatibility pair local models for $Mr.1$ with a ball on the left and models where this is not the case, and analogously for $Mr.2$.

We can now say what it means for a model to *satisfy* a formula of a language L_i. Let $\models_{cl}$ be the (classical) satisfiability relation between local models and formulas of L_i. Let us call $\models_{cl}$ *local satisfiability*. Notationally, let us write $i\colon\phi$ to mean ϕ and that ϕ is a formula of L_i. We say that ϕ is an L_i-formula, and that $i\colon\phi$ 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. Then we have the following:

Definition 3.2 (Satisfiability) *Let* $\mathbf{C}$ *be a model and* $i\colon\phi$ *a formula. Then,* $\mathbf{C}$ satisfies $i\colon\phi$, *in symbols* $\mathbf{C} \models i\colon\phi$ *iff, for all* $\mathbf{c} \in \mathbf{C}$, $\mathbf{c}_i \models i\colon\phi$, *where* $\mathbf{c}_i \models i\colon\phi$ *iff, for all* $m \in \mathbf{c}_i$, $m \models_{cl} \phi$.

Intuitively: a formula of the context C_i is satisfied by a model $\mathbf{C}$ if all the local models in C_i satisfy it. A model $\mathbf{C}$ satisfies a set of formulas Γ, in symbols $\mathbf{C} \models \Gamma$, iff $\mathbf{C}$ satisfies every formula $i\colon\phi \in \Gamma$.

The notions of model, context and satisfiability given above formalize the principles of locality and compatibility in the following sense:

Locality: Everything is local. First of all, the language is local: not only do we have a language for each context, but, also, there is no notion of a not labelled L_i-formula ϕ being satisfiable. We always talk of satisfiability of formulas in context, i.e., of labelled L_i-formulas. Second, 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. Third, the structures we consider to test local satisfiability are local: contexts have their own, generally different, domains of interpretation, sets of relations, and sets of functions.

Compatibility: Via compatibility, 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. On one extreme the two contexts have two independent views of the world. In this case we have a compatibility relation which is the empty set and there is no compatibility between what holds in the distinct sets of local models. On the other extreme the two contexts describe the same world. In this case all the languages are the same, and for every local model in a context there is a corresponding compatible identical local model in the other context. Furthermore, the compatibility relation is the identity mapping. In this case all the contexts are a replication of the same context and we are in the classical situation of one language and one notion of satisfiability and truth.

The notion of *validity* is the obvious one.

Definition 3.3 (Validity) *A formula $i\colon\phi$ is* valid, *in symbols $\models i\colon\phi$, if and only if all models satisfy $i\colon\phi$.*

What is more interesting is the notion of *logical consequence* which must take into account the fact that assumptions and conclusion may belong to distinct languages. Let Γ_j be a set of labelled L_j-formulas, and $\mathbf{c}_j$ be an element of a compatibility pair. Notationally, we write $\mathbf{c}_j \models \Gamma_j$ to mean that $\mathbf{c}_j \models j\colon\gamma$ for every $j\colon\gamma \in \Gamma_j$.

Definition 3.4 (Logical Consequence) *A formula $i\colon\phi$ is a* logical consequence *of a set of formulas $\Gamma = \cup_{j=1,2}\Gamma_j$, in symbols $\Gamma \models i\colon\phi$, if and only if, for all models $\mathbf{C}$, for all $\mathbf{c} \in \mathbf{C}$, if for all $j \in \{1,2\}$, $j \neq i$, $\mathbf{c}_j \models \Gamma_j$, then for all $m \in \mathbf{c}_i$, if $m \models_{cl} \Gamma_i$ then $m \models_{cl} \phi$*

Intuitively: take a formula $i\colon\phi$. Take a set of assumptions Γ and, among them, isolate the set of assumptions Γ_j with $j \neq i$. Take all the pairs whose local models in C_j satisfy Γ_j (and throw away all the others). Consider now the local models in C_i of the remaining pairs. $\Gamma \models i\colon\phi$ if in these remaining local models all the local models which satisfy Γ_i locally satisfy ϕ. Essentially, the idea is that the formulas in Γ_j prune away compatibility pairs, while the formulas in Γ_i prune away local models in C_i.

Notice that the notion of logical consequence defined above reduces to the classical notion of logical consequence in the case of a single language.

4 The two examples – formalized

4.1 Viewpoints

Let us start by defining L_1 and L_2. Let $P_1 = \{r, l\}$ and $P_2 = \{r, c, l\}$ be two sets of propositional constants (where, intuitively, r, c, l stand for ball on the right, in the center and on the left, respectively). Then L_1 and L_2 are defined as the smallest set containing P_1 and closed under the propositional connectives, and the smallest set containing P_2 and closed under the propositional connectives, respectively. Notice that L_1 (L_2) can describe all the views of $Mr.1$ ($Mr.2$) but that, for instance, L_1 cannot describe the situation with a ball in the center (that is, a possible view of $Mr.2$).

L_1 and L_2 have the usual Tarskian propositional semantics. The local models of L_1 are (univocally defined by the following sets of formulas):

$$\begin{array}{ll} m_1 = \{\neg l, \neg r\} & m_2 = \{l, \neg r\} \\ m_3 = \{\neg l, r\} & m_4 = \{l, r\}. \end{array}$$

The local models of L_2 are:

$$\begin{array}{ll} m_1 = \{\neg l, \neg c, \neg r\} & m_2 = \{l, \neg c, \neg r\} \\ m_3 = \{\neg l, c, \neg r\} & m_4 = \{l, c, \neg r\} \\ m_5 = \{\neg l, \neg c, r\} & m_6 = \{l, \neg c, r\} \\ m_7 = \{\neg l, c, r\} & m_8 = \{l, c, r\}. \end{array}$$

Let us construct a model for the magic box by imposing the following compatibility constraints:

1. if $Mr.1$ sees at least a ball then $Mr.2$ sees at least a ball;
2. if $Mr.2$ sees at least a ball then $Mr.1$ sees at least a ball;
3. $Mr.1$ and $Mr.2$ both see a box, or neither does.

The key notion in defining a compatibility relation is the identification of the set of L_i-formulas which are satisfied in all the local models in $\mathbf{c}_i$. Let us write $\Theta(\mathbf{c}_i)$ to mean such a set. $\Theta(\mathbf{c}_i)$, with $i = 1, 2$, characterizes the kind of compatibility imposed by a pair.

Definition 4.1 (A model for the magic box) *A model* $\mathbf{C}$ *for the magic box is a compatibility relation such that, for all* $\mathbf{c} \in \mathbf{C}$

$$\textit{if } l \vee r \in \Theta(\mathbf{c}_1) \textit{ then } l \vee c \vee r \in \Theta(\mathbf{c}_2); \tag{1}$$

$$\textit{if } l \vee c \vee r \in \Theta(\mathbf{c}_2) \textit{ then } l \vee r \in \Theta(\mathbf{c}_1); \tag{2}$$

$$\mathbf{c}_1 \neq \emptyset \textit{ iff } \mathbf{c}_2 \neq \emptyset. \tag{3}$$

Equation (1) models constraint 1. In fact, if $Mr.1$ sees a ball then this ball can only be on the left or on the right. In either case $l \vee r \in \mathbf{c}_1$, for all $\mathbf{c}$. Furthermore, in this case $Mr.2$ sees a ball in one of the three possible positions, and, therefore, $l \vee c \vee r \in \Theta(\mathbf{c}_2)$. A similar explanation can be given for Equation (2). Constraint 3 is more interesting. Equation (3) says that $Mr.1$ and $Mr.2$ have a model of the box only if they see it, and that this must hold for both of them, or for neither.

Example 4.1 Many models satisfying Definition 4.1 can be constructed. Let us consider the model(s) whose pairs only contain at most singleton sets. It is easy to see that in these models, if $Mr.1$ sees no balls then $Mr.2$ sees no balls, formally, $1{:}\neg l\wedge\neg r \models 2{:}\neg l\wedge\neg c\wedge\neg r$.

To prove this, let us consider all the pairs $\langle \mathbf{c}_1, \mathbf{c}_2\rangle$ such that $\mathbf{c}_1$ satisfies $1: \neg l \wedge \neg r$. From Definition 3.4 we must show that every $\mathbf{c}_2$ satisfies $2{:}\neg l \wedge \neg c \wedge \neg r$. Suppose that there exists a $\mathbf{c}_2$ such that $\mathbf{c}_2$ does not satisfy $2: \neg l \wedge \neg c \wedge \neg r$. $\mathbf{c}_2$ contains (at most) a propositional local model. Therefore, $\mathbf{c}_2$ satisfies $2: l \vee c \vee r$ and $l \vee c \vee r \in \Theta(\mathbf{c}_2)$. From Equation (1) we obtain $l \vee r \in \Theta(\mathbf{c}_1)$. $\mathbf{c}_1$ is a propositional model and it cannot satisfy both $1{:}\neg l \wedge \neg r$ and $1{:}l \vee r$. Thus $\mathbf{c}_1 = \emptyset$. By Equation (3), $\mathbf{c}_2 = \emptyset$ as well. This cannot be possible for Definition 3.1. Therefore the hypothesis that $\mathbf{c}_2$ does not satisfy $2{:}\neg l \wedge \neg c \wedge \neg r$ must be false. Thus $1{:}\neg l \wedge \neg r \models 2{:}\neg l \wedge \neg c \wedge \neg r$.

In a similar way, we can also prove the vice versa, that is, $2{:}\neg l\wedge\neg c\wedge\neg r \models 1{:}\neg l\wedge\neg r$. Furthermore, it can be shown that, in Definition 4.1, Equation (3) can be substituted with the following equation

$$\neg l\wedge\neg r \in \Theta(\mathbf{c}_1) \text{ iff } \neg l\wedge\neg c\wedge\neg r \in \Theta(\mathbf{c}_2) \tag{4}$$

which states exactly the two facts proved above. It is easy to notice that Equations (1), (2), (4) directly capture, as pairs with only singleton sets, all the compatible views in Figure 2 (Equation (4) capturing the one at the top). As a consequence the model which contains *all* the possible pairs containing only singleton sets is the model which directly formalizes the compatible views in Figure 2. The pairs in this model are as follows:

$$\left\{ \begin{array}{l} \langle\{\{\neg l, \neg r\}\}, \{\{\neg l, \neg c, \neg r\}\}\rangle \\ \langle\{\{l, \neg r\}\}, \{\{l, \neg c, \neg r\}\}\rangle, \\ \langle\{\{l, \neg r\}\}, \{\{\neg l, c, \neg r\}\}\rangle, \\ \ldots, \\ \langle\{\{l, r\}\}, \{\{l, c, r\}\}\rangle \end{array} \right\}$$

All the models whose pairs contain only singleton sets are subsets of this model.

4.2 Reasoning about belief

Let us start by defining L_a and L_b. L_a and L_b are two first order languages. L_a contains a predicate B and a term "ϕ" for every formula $\phi \in L_b$. $B(\text{“}\phi\text{”})$ intuitively stands for "b believes ϕ". To make things more concrete, let us suppose that b has beliefs about a young student, named Paul. To this extent, let us add an individual constant p (which stands for Paul) and two unary predicates S and Y (intuitively meaning "to be student" and "to be young", respectively) to L_b.

L_a and L_b have the usual Tarskian semantics. The local models of L_a are pairs $\langle \mathbf{dom}_a, \mathcal{I}_a\rangle$, where: $\mathbf{dom}_a$ is the domain of interpretation of a, and $\mathcal{I}_a$ maps every term "ϕ" to a constant $d_\phi \in \mathbf{dom}_a$, and the predicate B to a relation $R_B \subseteq \mathbf{dom}_a$. The local models of L_b are pairs $\langle \mathbf{dom}_b, \mathcal{I}_b\rangle$ where: $\mathbf{dom}_b$ is the domain of interpretation of b, and $\mathcal{I}_b$ is an interpretation function mapping the individual constant p to a constant $d_p \in \mathbf{dom}_b$ and the predicates S and Y to the relations $R_S \subseteq \mathbf{dom}_b$ and $R_Y \subseteq \mathbf{dom}_b$, respectively.

The key part in the construction of models of reasoning about the belief is the definition of the relation existing between the beliefs of b and the beliefs of a of the form $B(\text{“}\phi\text{”})$. To this extent let us introduce the following notation:

- Let Γ be a set of L_b-formulas. We write $B(\text{“}\Gamma\text{”})$ to mean the set of L_a-formulas $B(\text{“}\phi\text{”})$ such that ϕ belongs to Γ.
- Let Γ be a set of L_a-formulas. We write $B^{-1}(\text{“}\Gamma\text{”})$ to mean the set of L_b-formulas ϕ such that $B(\text{“}\phi\text{”})$ belongs to Γ.

Let us define a model with a correct and complete with respect to b's beliefs. In this case "b believes ϕ" belongs to a's beliefs iff ϕ belongs to b's beliefs.

Definition 4.2 (A Model for reasoning about belief) *A model* $\mathbf{C}$ *for reasoning about belief is a compatibility relation such that, for all* $\mathbf{c} \in \mathbf{C}$

$$B^{-1}(\text{“}\Theta(\mathbf{c}_a)\text{”}) \subseteq \Theta(\mathbf{c}_b) \tag{5}$$

$$B(\text{“}\Theta(\mathbf{c}_b)\text{”}) \subseteq \Theta(\mathbf{c}_a) \tag{6}$$

(5) models the soundness of a; (6) models its completeness. Let us explain (6) ((5) can be explained analogously). Take $\Theta(\mathbf{c}_b)$, that is all the L_b-formulas which are true in $\mathbf{c}_b$. $\Theta(\mathbf{c}_b)$ is the set of the beliefs of b, at least for what concerns $\mathbf{c}$ (taking a belief to be a formula which is true in all local models). Take now the set of a's beliefs about the actual beliefs of b, that is $B(\text{“}\Theta(\mathbf{c}_b)\text{”})$ (again restricted to $\mathbf{c}$). Then for a to be complete (but not necessarily correct) we must have that $B(\text{“}\Theta(\mathbf{c}_b)\text{”})$ is a subset of its beliefs, in symbols $B(\text{“}\Theta(\mathbf{c}_b)\text{”}) \subseteq \Theta(\mathbf{c}_a)$.

Example 4.2 Suppose that a believes that b believes that Paul is a student and suppose that b believes that every student is young. Does a believe that b believes that Paul is young? Formally, we want to verify whether $a\colon B(\text{“}Y(p)\text{”})$ is a consequence of the two axioms $a\colon B(\text{“}S(p)\text{”})$ and $b\colon \forall x.S(x) \supset Y(x)$.

All models $\mathbf{C}$ (and all $\mathbf{c}_i$) satisfy $a\colon B(\text{“}S(p)\text{”})$ (being this formula an axiom). From the definition of satisfiability $B(\text{“}S(p)\text{”}) \in \Theta(\mathbf{c}_a)$ for every $\mathbf{c}_a$. From Equation (5), $S(p) \in \Theta(\mathbf{c}_b)$ for every $\mathbf{c}_b$ and thus $\mathbf{C}$ satisfies $b\colon S(p)$. All models $\mathbf{C}$ satisfy $b\colon \forall x.S(x) \supset Y(x)$. From the definition of satisfiability, every local model of L_b in $\mathbf{C}$ satisfies $\forall x.S(x) \supset Y(x)$. Local models are classical models, thus they satisfy $S(p) \supset Y(p)$ and, therefore, $Y(p)$. It follows that $\mathbf{C}$ satisfies $b\colon Y(p)$, i.e. $b\colon Y(p) \in \Theta(\mathbf{c}_b)$. By applying Equation (6) we obtain $a\colon B(\text{“}Y(p)\text{”}) \in \Theta(\mathbf{c}_a)$, i.e. the model satisfies $a\colon B(\text{“}Y(p)\text{”})$. This completes the proof.

It is important to notice that B is a predicate. Its extension are the formulas it holds on. It is natural to ask what is the relation between the semantics proposed here and other approaches, e.g., possible worlds semantics, where belief is formalized as a modal operator. Section 5 provides an informal but detailed comparison between Local Models Semantics and possible worlds semantics in general. Limited to belief, we start by observing that the extension of B and the beliefs of b are defined independently, and that Equations (5) and (6) only provide a posteriori compatibility constraints. This feature is (in part, see also Section 5) the semantic counterpart of the modularity and incrementality hinted in Section 2 and discussed at the proof-theoretical level in [Giunchiglia *et al.*, 1993]. It allows us to define the kind (and quality) of a's beliefs about beliefs by playing with equations like Equations (5) and (6). For instance, it can be shown that Local Models Semantics scales up very naturally from very strong forms of non-omniscient belief up to omniscient belief, and that it allows for a very elegant treatment of bounded beliefs. Technically, it is possible to prove completeness and correctness results with the proof calculi for belief (modeled as contextual reasoning) described in [Giunchiglia and Serafini, 1994; Giunchiglia *et al.*, 1993]. (These proofs are not given here for lack of space.). The proof calculi presented in these two papers (and others) can be proved equivalent with the various normal and non normal modal logics, e.g., those described in [Fagin and Halpern, 1988]. All these results sound as follows. Take a modal logic X_M and the corresponding logic for contextual reasoning X_C. Then

$$\vdash_{X_M} \phi \text{ iff } \vdash_{X_C} a\colon \phi.$$

In other words, a (or, in the case of a hierarchy of agents, the agent at the top of the hierarchy) proves the same set of theorems as the modal logic under consideration. The logic for reasoning about belief defined in this section allows for no nesting of B, and makes B weaker than the modal operator of modal K.

5 Other frameworks – a comparison

The obvious, most studied, framework to start from is possible worlds semantics. Both Local Models Semantics and possible worlds semantics allow for multiple objects (models or worlds) and have a notion of local satisfiability (to a local model or to a possible world). However there are also some important differences. First, in possible worlds there is a unique language which describes what is true in all the worlds and there is no notion of truth of a labelled formula. This is the case also for the extensions of possible worlds semantics aimed at formalizing local reasoning (see, e.g., [Fagin and Halpern, 1988]), where localization is achieved by adding a new modal operator to the language. Second, worlds are not (Tarskian) models, the key difference being that possible worlds allow for the use of modal operators. The satisfiability of a formula containing a modal operator is defined in terms of the accessibility relation, which must therefore be given before defining satisfiability in a world. The notion of satisfiability in a world is a function of the model of which the world is part. This is not the case for Local Models Semantics where each local model has its own notion of satisfiability, and where the model and its structure influence only the set of local models under consideration. The hypothesis of using a single unique global language and of being able to describe a priori the structure of the model under consideration is very useful and works in many situations. It does not seem to work in those cases where there is no global scheme describing the system, e.g., the federation of heterogeneous data or knowledge bases or multiagent systems.

In the last few years various semantics for contextual reasoning have been proposed. Most of them are based on possible worlds semantics. As far as we know, the first attempt is described in [Giunchiglia *et al.*, 1992]. In this work there is a notion of labelled formula and of (local) satisfiability to a set of (local) possible worlds. This semantics works well for contextual logics equivalent to modal K or stronger. Its main limitation is that it is not clear how to extend it to other logics, e.g., non normal modal logics or logics for reasoning with viewpoints.

Guha, in his PhD thesis [Guha, 1991] informally describes a semantics for reasoning with context. Understanding Guha's informal definitions is a non-trivial task. Some of the main ideas seem the following. There is a single global language from which it is possible to extract the (local) languages of all the contexts. There seems to be a notion of satisfaction of labelled formulas, and a notion of labelled formulas being meaningless in a context. There is distinguished symbol `ist`, whose intuitive meaning is "is true", which seems treated as a modal operator. Guha's semantics has been partially formalized in the work by Buvac and his co-authors (see for instance [Buvac and Mason, 1993]). Buvac's semantics seems to have the same features and defects as the semantics in [Giunchiglia *et al.*, 1992], with the further complication that, starting from a single language, there is a lot of work to do in order to achieve locality. In particular the formulas of the global language which are meaningless in a context must be treated as such (this is done using Bochvar three valued logic).

6 Conclusion

In this paper we have presented a new semantics, called Local Models Semantics, and proposed it as a foundation for reasoning with context. Local Models Semantics formalizes the two general principles underlying contextual reasoning, namely the principle of locality and the principle of compatibility. In this paper we have shown how Local Models Semantics can be used to model two important forms of contextual reasoning, namely reasoning with viewpoints and reasoning about belief.

References

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

[Buvac and Mason, 1993] Sasa. Buvac and Ian A. Mason. Propositional logic of context. *Proc. of the 11th National Conference on Artificial Intelligence*, 1993.

[Fagin and Halpern, 1988] R. Fagin and J.Y. Halpern. Belief, awareness, and limited reasoning. *Artificial Intelligence*, 34:39–76, 1988.

[Farquhar *et al.*, 1995] 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.

[Fauconnier, 1985] G. Fauconnier. *Mental Spaces: aspects of meaning construction in natural language*. MIT Press, 1985.

[Giunchiglia and Serafini, 1994] F. Giunchiglia and L. Serafini. Multilanguage hierarchical logics (or: how we can do without modal logics). *Artificial Intelligence*, 65:29–70, 1994. Also IRST-Technical Report 9110-07, IRST, Trento, Italy.

[Giunchiglia *et al.*, 1992] F. Giunchiglia, L. Serafini, and A. Simpson. Hierarchical meta-logics: intuitions, proof theory and semantics. In *Proc. of META-92, Workshop on Metaprogramming in Logic*, number 649 in LNCS, pages 235–249, Uppsala, Sweden, 1992. Springer Verlag. Also IRST-Technical Report 9101-05, IRST, Trento, Italy.

[Giunchiglia *et al.*, 1993] 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. Also IRST-Technical Report 9206-03, IRST, Trento, Italy.

[Giunchiglia, 1993] 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. Also IRST-Technical Report 9211-20, IRST, Trento, Italy.

[Guha, 1991] 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.

[Kaplan, 1978] D. Kaplan. On the Logic of Demonstratives. *Journal of Philosophical Logic*, 8:81–98, 1978.

[McCarthy, 1987] 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.

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

[Serafini and Ghidini, 1997] L. Serafini and C. Ghidini. Context Based Semantics for Federated Databases. In *Proceedings of the 1st International and Interdisciplinary Conference on Modeling and Using Context (CONTEXT-97)*, pages 33–45, Rio de Jeneiro, Brazil, 1997. Also IRST-Technical Report 9609-02, IRST, Trento, Italy.