Updating Contexts*

Antonia Donà and Luciano Serafini
ITC – IRST
38050 Povo, Trento, Italy
{antodona, serafini}@itc.it

Abstract

Different formalisms for contexts have been successfully used to represent and reason about distributed knowledge. In these formalisms, knowledge is represented as a set of contexts, each representing a piece of the whole knowledge. Contexts have been proved particularly adapt to deal with heterogeneous distributed knowledge, but, in order to deal with the change of such a knowledge, they must be extended. In particular, they must cope with the problems of adding, deleting, or changing facts in a context, and computing the effect of this change in the other contexts. The current approaches to belief revision and multi-agent belief revision can be helpful, but they do not provide a satisfactory treatment of heterogeneity. We provide a formal definition of the operation of updating a context, and we define an algorithm that computes the effects of updating a context on other related contexts. We take a semantic perspective, i.e., each context is formalized as a set of possible partial models of the world.

*We would like to thank Marco Aiello, Roberta Ferrario, Fausto Giunchiglia, and Ilya Zaihrayeu for useful feedback and proof readings, This work is partially supported by the project Trentino Experience funded by the Provincia Autonoma di Trento.

1 Introduction

In a distributed environment, such as the information system of a large company, or the beliefs of a set of agents, or, to the extreme, the whole world wide web, knowledge is not represented as a monolithic system. Rather it is scattered in a number of relatively “small” and specialized pieces, called local knowledge bases. Each local knowledge base adopts a language to represent its knowledge. This language can either be shared among all, or a subset of the local knowledge bases, or it could be a specialized language. Furthermore, each local knowledge base gives a way to interpret (i.e. it associates with a semantics) the statements of its language in terms of real world entities (i.e. propositions or objects). In distributed knowledge we cannot assume that this semantics is the same for all the local knowledge bases, but the well known phenomenon (and problem) of semantic heterogeneity arises. Semantic heterogeneity occurs when the same or related real world entities can be represented, at a different level of abstraction and by different statements, in more than one local knowledge base. This means that local knowledge bases are not independent, rather they overlap.

Due to overlapping, changing an information of a knowledge base results both into an update of the local knowledge base and into a revision of the knowledge in those bases which overlap with the original one. These revisions might in turn lead to revising of other knowledge bases, and so on. Determining the minimal update propagation in order to maintain consistency in a distributed knowledge base has recently become a very relevant problem. The main approaches in this area are those in [19, 11, 14, 12, 13]. These approaches, however, do not provide a satisfactory solution for revising distributed knowledge (or beliefs) in presence of heterogeneity. The main reason is that all rely on the fact that shared knowledge, i.e. the overlaps between local knowledge bases, has a uniform representation. This assumption, however, in many cases, is not realistic. Consider the following example.

Example 1. Consider the scenario in which $Mr.1$ and $Mr.2$ are two agents looking at a box composed of two slots, each possibly containing a ball (see the Figure 1). Suppose that $Mr.1$ and $Mr.2$ can move around the box along two circular tracks, and that $Mr.1$ moves on the internal track, while $Mr.2$ moves on the external one. According to what they can see, $Mr.1$'s beliefs are about balls being in the right or left slot of the box, and $Mr.2$'s beliefs are about balls being in the left/right slot, and about a man being between himself and the box. For instance, in the situation shown on the left side of Figure 1, $Mr.1$ believes that there is a ball on the left slot, and that there is no ball on the right slot (in symbols $L \wedge \neg R$).

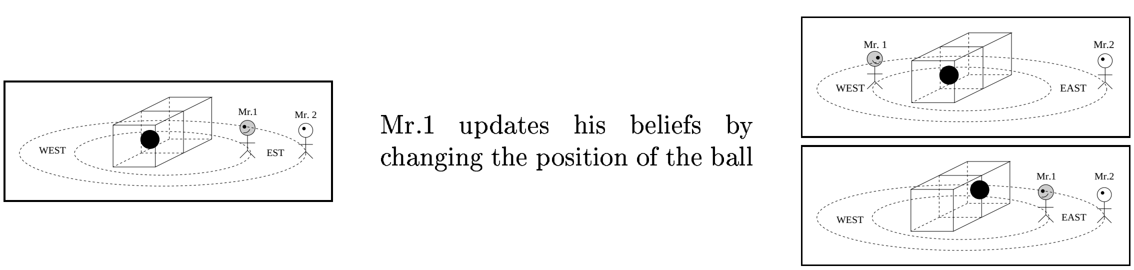

Initial beliefs

Possible updated beliefs

Figure 1: *An example.*

$Mr.2$ believes that there is a ball on the left slot, that there is no ball on the right slot, and that there is a man in front of him (in symbol $L \wedge \neg R \wedge M$).

Notice that in this example, the relation between the meaning of L in $Mr.1$ and $Mr.2$ beliefs depends on the value of M in $Mr.2$ beliefs. If M is true, then L has the same meaning in both $Mr.1$ and $Mr.2$'s beliefs. If M is false, the meaning of L in $Mr.1$'s beliefs is independent from the meaning of L in $Mr.2$'s beliefs as they are referring to two different slots.

Suppose that, after a while $Mr.1$ is looking at the box and he sees that the ball has been moved from the left side to the right side of the box. He therefore updates his beliefs to $\neg L \wedge R$. Suppose also that $Mr.1$ communicates to $Mr.2$ his new beliefs. How does $Mr.2$ react to this, if he trusts $Mr.1$? He has two options, either to believe that $Mr.1$ has changed his position going to other side of the box, or that the ball has been moved from the left to the right. In the first case $Mr.1$ updates its beliefs to $L \wedge \neg R \wedge \neg M$, in the second case his beliefs are updated to $\neg L \wedge R \wedge M$. These two updates correspond to the two scenes shown on the right of Figure 1.

A formalization of distributed belief revision should provide a solution to the following problem: consider the implementation of two autonomous computer programs (agents) that simulate $Mr.1$ and $Mr.2$ reasoning, without having access to a description of the state of the world (in terms of $Mr.1$ and $Mr.2$ positions and the positions of the balls in the box), but simply relying on the coordination between the two agents. If, for any reason one agent decides to revise its beliefs, how does the other agent revise its own beliefs to maintain them consistent with a possible (not accessible) state of the world?

In this paper, we propose a new logic based approach to the revision of distributed knowledge assuming there is no global representation of the knowledge, and that information in different knowledge bases can be related (e.g., the same information is represented in two or more knowledge bases, the information in a knowledge base is constrained to the information in a different knowledge base).

Our approach is based on the theory of context. Contexts formalisms have been proposed in AI by McCarthy, Buvač and Mason [3, 4], and by Giunchiglia, Ghidini, and Serafini [8, 6, 18]. Furthermore, in [17, 15] *contexts* have been proposed for the representation of distributed knowledge in presence of heterogeneity. The theory so far, however, does not provide any support for revising the contents of contexts. In this paper, we give a formal definition of contextual update, and we provide a sound and complete procedure for computing the possible outcomes of a contextual update.

The structure of the paper is the following. In Section 2, we introduce the formalisms, based on contexts, which we have adopted. In Section 3, we give a set of criteria on the update of context, which lead us to an implicit definition of the update operation. In Section 4, we provide an algorithm that computes this update operator. We conclude the paper by a comparison with similar approaches (Section 5).

2 Logic of contexts

To represent distributed knowledge we adopt the logic of contexts developed by Giunchiglia, Serafini, and Ghidini [8, 6]. In this section we recall the key concepts of this logic.

Logic of context represents knowledge in a set of theories (contexts) each of them constituting a partial, approximate, and perspectival description of the world [7, 2]. Each context partially describes a portion of the world from a perspective. For instance, the context of $Mr.1$'s beliefs partially describes the content of the box, as $Mr.1$ might not believe where a ball is. It describes only a portion of the world, as in $Mr.1$'s context there are no beliefs about the position of $Mr.2$; finally, it describes the world from $Mr.1$'s contextual perspective, because the meaning of the formulas in this context depends on $Mr.1$'s position.

To each context we associate a logical language that, for the purpose of this paper, is supposed to be a propositional language. Formally let $\{L_i\}_{i\in I}$ be a family of propositional languages defined over a set of indices I (in the following simply $\{L_i\}$). Intuitively, I contains the names of the contexts in which knowledge (beliefs) is partitioned and each L_i is the content language of the i-th context.

In our example, we can define the propositional languages L_1 and L_2 used by $Mr.1$ and $Mr.2$ to describe their views. L_1 is the propositional language with the primitive propositions $\{L, R\}$ and L_2 is the propositional language with the primitive propositions $\{L, R, M\}$.

The meaning of formulas depends on the context where they are stated. For instance, the meaning of L in context 1 is that $Mr.1$ believes that there is a ball on the left side of the box. But the left depends on the position of $Mr.1$. The proposition M in the context 1 has no meaning at all. To distinguish the different meaning of propositions, let us write $i:\phi$ to mean ϕ and that ϕ is a formula of L_i. We say that ϕ is an L_i-formula, and that $i:\phi$ is a formula.

The semantics of a family of languages $\{L_i\}$, called Local Model Semantics, is based on the notion of chain. Chains are defined in terms of propositional models of the languages $\{L_i\}_{i\in I}$. A propositional model m of L_i is a consistent and complete set of literals in L_i, i.e,, a set m such that, for any primitive proposition p of L_i, either $p \in m$ or $\neg p \in m$, but not $\{p, \neg p\} \subseteq m$.

Definition 1 (Chain). A *chain* is a function c that associates to each index $i \in I$ a set of models of L_i, such that there is at least an i with $c(i) \neq \emptyset$. We write c_i instead of $c(i)$. A chain c satisfies a formula $i:\phi$, in symbols $c \models i:\phi$, if for any $m \in c_i$, $m \models \phi$, according to the definition of satisfiability ($\models$) in propositional logic.

Intuitively, a chain represents a combination of epistemic states, one for each context. The epistemic state associated to each context is a set of possible states of the "piece of world" described from the perspective of the context. Two chains in the languages for Example 1 are:

$$c = \left\langle \boxed{\bullet}\boxed{}, \boxed{\bullet}\boxed{}\boxed{\text{M}} \right\rangle \quad (1)$$

$$c' = \left\langle \boxed{\bullet}\boxed{\bullet}, \left\{ \boxed{\bullet}\boxed{}\boxed{\text{M}}, \boxed{}\boxed{\bullet}\boxed{\text{M}} \right\} \right\rangle \quad (2)$$

The graphical notation $\boxed{\bullet}\boxed{}$ represents the model $\{L, \neg R\}$ of the language L_1. Similarly, $\boxed{\bullet}\boxed{}\boxed{\text{M}}$ represents the model $\{L, \neg R, M\}$ of L_2. The other graphical notations have analogous meaning. Furthermore, if a set is composed of a single element we omit the curly brackets.

The chain c represents the (possible) situation where $Mr.1$ and $Mr.2$ both believe that there is ball on the left side and no ball on the right side and $Mr.2$ believes also that there is a man between himself and the box. This combination of epistemic states is possible as it corresponds to the situation shown on the left box of Figure 1). Notice, however, that c is also consistent with the situation where $Mr.1$ and $Mr.2$ are both on the west side of the box, and the ball is moved to the other slot. The local perspectives of $Mr.1$ and $Mr.2$, however, cannot discriminate between these two situations. c' represents a situation where $Mr.1$ believes that there are two balls in the box, while $Mr.2$ believes that there is only one ball. This is an *impossible* combination of epistemic states, as there is no real situation where this can happen.

Definition 2 (Compatibility relation). A *compatibility relation* is a non empty set C of chains. A compatibility relation C satisfies $i : \phi$, in symbols $C \models i : \phi$, if, for each chain $c \in C$, $c \models i : \phi$.

Intuitively, a compatibility relation contains the combinations of epistemic states compatible with a "real situation". For instance, the compatibility relation associated with Example 1 contains the chain c and does not contain chain c' defined in (1) and (2) respectively. c is indeed compatible with the situation shown on the left box of Figure 1 while c' is not compatible with any possible state of the world of this example. The fact that a chain actually corresponds to a real situation, however, is not explicitly represented in the semantics, as a compatibility relation does not contain any formalization of the state of the world.

Despite the fact that a compatibility relation can be *any* set of chains, practical use, restricts us to consider compatibility relations that can be finitely specified using some specification language. From the theory of context defined in [8] we take the notion of *bridge rules*. Bridge rules allow to relate the truth of a formula in a context i with the truth of a set of formulas in other contexts.

Definition 3 (Bridge rule). Let $i_1, \ldots, i_n, j \in I$, $i_1 : \phi_1, \ldots, i_n : \phi_n$ and $j : \psi$ be propositional formulas. A *bridge rule* is an expression of the form: $i_1 : \phi_1, \ldots, i_n : \phi_n \rightarrow j : \psi$. A chain c satisfies a bridge rule in case, if $c \models i_1 : \phi_1, \ldots, c \models i_n : \phi_n$, then $c \models j : \psi$.

A set of bridge rules BR defines the compatibility relation C_{BR} composed of all the chains that satisfy BR. Logical consequence in these compatibility relations captures the relations among formulas in different contexts, modeled by a compatibility relation.

Definition 4 (Logical Consequence w.r.t. a set of BR). A formula $i : \phi$ is a *logical consequence* of a set of formulas Γ w.r.t. a set of bridge rules BR ($\Gamma \models_{BR} i : \phi$) if, for all $c \in C_{BR}$, if $c \models \Gamma$, then $c \models i : \phi$

Example 2. A set of bridge rules that formalize the constraints between the beliefs of $Mr.1$. and $Mr.2$ for Example 1 are shown in Figure 2.

Bridge rules have a declarative meaning, which is the one provided by the definition of satisfiability (Definition 3). They also have a procedural meaning in terms of update propagation. For instance, the bridge rules BR1, BR2, BR5 and BR6 state that updates on L in 1 and M in 2 affect the update of L and R in 2. Similar interpretations hold for all the other bridge rules. Notice that, with this intuitive interpretation, M cannot be changed as an effect of an update on L and R, as M does not appear in the consequence of any bridge rule.

3 The contextual update operator

Example 3. Consider once again Example 1. Suppose that both $Mr.1$ and $Mr.2$ believe that there is a ball on the left side and no balls on the right side of the box, and that $Mr.2$ believes that $Mr.1$ is in front of him. The chain representing this beliefs of $Mr.1$ and $Mr.2$ is the following:

$$c = \left\langle c_1 = \begin{array}{|c|c|}\hline \bullet & \\ \hline \end{array},\ c_2 = \begin{array}{|c|c|c|}\hline \bullet & & \mathrm{M} \\ \hline \end{array} \right\rangle \tag{3}$$

Now suppose that $Mr.1$ updates his beliefs based on the fact that he perceives that the ball moves from the left to the right side of the box. The chain formalizing this situation is:

$$c' = \left\langle c'_1 = \begin{array}{|c|c|}\hline & \bullet \\ \hline \end{array},\ c'_2 = \begin{array}{|c|c|c|}\hline \bullet & & \mathrm{M} \\ \hline \end{array} \right\rangle \tag{4}$$

This chain does not satisfy bridge rules BR3, BR5, BR9, and BR15. Indeed it describes an impossible situation. We need to minimally update $Mr.2$, in order to obtain a chain that satisfies all the bridge rules. There are the following two possibilities:

$$c'' = \left\langle c''_1 = \begin{array}{|c|c|}\hline & \bullet \\ \hline \end{array},\ c''_2 = \begin{array}{|c|c|c|}\hline \bullet & & \\ \hline \end{array} \right\rangle \tag{5}$$

$$c''' = \left\langle c'''_1 = \begin{array}{|c|c|}\hline & \bullet \\ \hline \end{array},\ c'''_2 = \begin{array}{|c|c|c|}\hline & \bullet & \mathrm{M} \\ \hline \end{array} \right\rangle \tag{6}$$

c'', however, is not acceptable as it involves a change of M, without $Mr.2$ having decided to change explicitly. We therefore select c'''. Finally, suppose that $Mr.2$ decides to update M obtaining the following chain:

$$c'''' = \left\langle c''''_1 = \begin{array}{|c|c|}\hline & \bullet \\ \hline \end{array},\ c''''_2 = \begin{array}{|c|c|c|}\hline & \bullet & \\ \hline \end{array} \right\rangle \tag{7}$$

Again c'''' does not satisfies the bridge rules. In this case, however, there is not a unique possibility. Possible changes to restores consistency are the following:

$$c^1 = \left\langle c^1_1 = \begin{array}{|c|c|}\hline & \bullet \\ \hline \end{array},\ c^1_2 = \begin{array}{|c|c|c|}\hline \bullet & & \\ \hline \end{array} \right\rangle \qquad c^3 = \left\langle c^3_1 = \begin{array}{|c|c|}\hline \bullet & \bullet \\ \hline \end{array},\ c^3_2 = \begin{array}{|c|c|c|}\hline \bullet & \bullet & \\ \hline \end{array} \right\rangle$$

$$c^2 = \left\langle c^2_1 = \begin{array}{|c|c|}\hline \bullet & \\ \hline \end{array},\ c^2_2 = \begin{array}{|c|c|c|}\hline & \bullet & \\ \hline \end{array} \right\rangle \qquad c^4 = \left\langle c^4_1 = \begin{array}{|c|c|}\hline & \\ \hline \end{array},\ c^4_2 = \begin{array}{|c|c|c|}\hline & & \\ \hline \end{array} \right\rangle$$

In terms of number of changes the above adjustments involve two changes each and therefore there is no theoretical a priori reason to prefer one to the other. By taking a skeptical approach, we leave open the possibility of all of them, by adopting a compatibility relation, instead of a single chain. Therefore we have that the new beliefs of $Mr.1$ and $Mr.2$ are formalized by the compatibility relation composed of the four chains.

Let us now give a formal definition of update so that the previous example can be represented in a systematic way.

(BR1)	$1{:}L, 2{:}M$	$\rightarrow$	$2{:}L$	(BR2)	$1{:}L, 2{:}\neg M$	$\rightarrow$	$2{:}R$
(BR3)	$1{:}R, 2{:}M$	$\rightarrow$	$2{:}R$	(BR4)	$1{:}R, 2{:}\neg M$	$\rightarrow$	$2{:}L$
(BR5)	$1{:}\neg L, 2{:}M$	$\rightarrow$	$2{:}\neg L$	(BR6)	$1{:}\neg L, 2{:}\neg M$	$\rightarrow$	$2{:}\neg R$
(BR7)	$1{:}\neg R, 2{:}M$	$\rightarrow$	$2{:}\neg R$	(BR8)	$1{:}\neg R, 2{:}\neg M$	$\rightarrow$	$2{:}\neg L$
(BR9)	$2{:}L \wedge M$	$\rightarrow$	$1{:}L$	(BR10)	$2{:}L \wedge \neg M$	$\rightarrow$	$1{:}R$
(BR11)	$2{:}R \wedge M$	$\rightarrow$	$1{:}R$	(BR12)	$2{:}R \wedge \neg M$	$\rightarrow$	$1{:}L$
(BR13)	$2{:}\neg L \wedge M$	$\rightarrow$	$1{:}\neg L$	(BR14)	$2{:}\neg L \wedge \neg M$	$\rightarrow$	$1{:}\neg R$
(BR15)	$2{:}\neg R \wedge M$	$\rightarrow$	$1{:}\neg R$	(BR16)	$2{:}\neg R \wedge \neg M$	$\rightarrow$	$1{:}\neg L$

Figure 2: Bridge rules among $Mr.1$ and $Mr.2$ beliefs.

3.1 Implicit definition

A contextual update operator takes a local update (an update in a context) a compatibility relation and returns a compatibility relation in which the local update is performed and all the bridge rules are satisfied. Local updates are specified in a *local update language*.

Definition 5 (Local update languages). Given a propositional language L a *local update* in L is defined as follows:

1. For any propositional letter p, $\mathsf{add}(p)$ and $\mathsf{del}(p)$ are atomic local updates.
2. If upd_1 and upd_2 are local updates, then $\mathsf{upd}_1; \mathsf{upd}_2$ is a local update.
3. Nothing else is a local update.

A local update has a *main effect*, explicitly specified in the operator, and a derived effect, not specified in the operator. The former is the set of local updates that are necessary to keep a chain consistent with respect to the set of bridge rules. Intuitively, the main effect of $\mathsf{add}(p)$ and $\mathsf{del}(p)$ is that p is set to true and false respectively; the main effect of $\mathsf{upd}_1; \mathsf{upd}_2$, is the effect of the sequential execution of upd_1 and upd_2. The following definition should clarify this intuition. In the following, l denotes a literal, i.e. a propositional letter, or the negation of a propositional letter. l^c denotes the complement of the literal l, i.e. p, if l is $\neg p$, and $\neg p$ if l is p.

Definition 6 (Main Effect). Let upd be a local update, the *main effect* of upd is defined as follows:

1. $me(\mathsf{add}(p)) = \{p\}$
2. $me(\mathsf{del}(p)) = \{\neg p\}$
3. $me(\mathsf{upd}_1; \mathsf{upd}_2) = me(\mathsf{upd}_2) \cup (me(\mathsf{upd}_1) \setminus me(\mathsf{upd}_2)^c)$, where for any set of literals S, $S^c = \{l^c | l \in S\}$.

Let's comment point 3 of the previous definition. The main effect of a sequence of two updates is the union of the effects of the two updates, minus the effect of the first update that are overwritten by the second one. For example, the effect of $me(\mathsf{del}(p); \mathsf{add}(p))$ is that p is true on the final state of the context, and from the given definition is equal to $\{p\} \cup (\{\neg p\} \setminus \{\neg p\}) = \{p\}$.

Let us now focus on the definition of the derived effects of a local update. For these effects we do not give an explicit definition, we rather give three *principles* and implicitly define the update operation as the one that satisfy the principles. The principles are:

Admissibility The result of an update must satisfy a given set of bridge rules.

Consistency The result of an update must be a consistent set of epistemic states.

Necessary changes Each derived effect should be "justifiable" to be necessary in order to maintain consistency. Here by justifiable we mean that the derived effect is a logical consequence of the main effect and the unchanged part.

The definition of the update operation is obtained by rephrasing the previous principles in a more formal way.

Definition 7 (Contextual update). Let upd be a local update on the language L_i and BR be a set of bridge rules. The *contextual update*, $i:\mathsf{upd}$ is a function that, for each chain $c \in C_{BR}$ returns a compatibility relation $i:\mathsf{upd}(c)$ containing all the chains c' such that:

1. c' satisfies the set BR of bridge rules;
2. $i:me(\mathsf{upd}), \Sigma(c, c') \models \Delta(c, c')$;
3. $i:me(\mathsf{upd}), \Sigma(c, c') \not\models i:\bot$.

where: $\Sigma(c, c')$, the common part of c and c', is equal to $\{i:\phi \mid c \models i:\phi \text{ and } c' \models i:\phi\}$, and $\Delta(c, c')$, the difference between c and c', is equal to $\{i:\phi \mid c \not\models i:\phi \text{ and } c' \models i:\phi\}$.

Example 4. Consider Example 3, with the initial situation formalized by chain (3). Suppose that $Mr.1$ updates his beliefs as described in the example by applying the local update $\mathsf{del}(L); \mathsf{add}(R)$. Consider the chains (5) and (6) both satisfy condition 1 and 3 of Definition 7 on the set of bridge rules BR1–16, but only (6) is acceptable as (5) does not satisfy condition 2 of Definition 7. Indeed, we have that the change of M into $\neg M$ in 2 is not justifiable in terms of logical consequence as $1:\neg L, 1:R, 2:\neg R \not\models 2:\neg M$.

Notice that there are cases where there is no c' that satisfies point 1-3 of Definition 7 and therefore $i:\mathsf{upd}(c)$ is not defined. In this case update is not allowed since it conflicts with one of the three principles.

Example 5. Let $c = \{\{p\}, \{q\}\}$ be a chain on two languages L_1 and L_2 containing the propositions $\{p\}$ and $\{q\}$ respectively. Consider the single bridge rule $1:p \rightarrow 2:q$. The set $2:\mathsf{del}(q)(c)$ is the empty set. Indeed if q is set to false in 2, in order to satisfy the rule, some change should be performed in 1. On the other hand, there is no way to infer such a derived change in 1. Notice the absence of bridge rules that go from 2 to 1.

3.2 Properties

In this section we compare the contextual update operator with two of the main references in belief update: the foundation paper of Katsuno and Mendelson [10], and the review paper of Herzig [9]. Contextual update satisfies the four basic desiderata for belief revision described in the [9].

Proposition 1. *If $i:\mathsf{upd}(c)$, is defined, then:*

1. *Consistency: $i:\mathsf{upd}(c) \not\models i:\bot$.*
2. *Syntax Independence: If $i:me(\mathsf{upd}_1) = i:me(\mathsf{upd}_2)$, and c and c' are equivalent (i.e. they satisfy the same formulas), then $i:\mathsf{upd}_1(c) = i:\mathsf{upd}_2(c')$ are equivalent.*
3. *Success: $i:\mathsf{upd}(c) \models i:me(\mathsf{upd})$.*
4. *Minimal changes: For all $c' \in i:\mathsf{upd}(c)$, then for all c'' such that $\Delta(c, c') \subset \Delta(c, c'')$, $c'' \notin i:\mathsf{upd}(c)$.*

Belief revision is usually defined for formulas, while here we consider only literals. In order to compare our approach with the traditional AGM postulates [1] of belief revision and the KM postulates [10] for belief update, we need to generalize our update operators to the case of complex formulas $i:\phi$.

For an i-formula ϕ, let $\text{DNF}_\phi = \{S_1, \ldots, S_k\}$ be the disjunctive normal form of ϕ, and let upd_ϕ be the set of update operations defined as follows: for any $S = \{p_1, \ldots, p_k, \neg q_1, \ldots, \neg q_k\} \in \text{DNF}_\phi$, upd_ϕ contains the update operation $\mathsf{upd}_S = \mathsf{add}(p_1); \ldots; \mathsf{add}(p_n); \mathsf{del}(q_1); \ldots; \mathsf{del}(q_n)$. The generalized update operation $C \circ i\!:\!\phi$ for a compatibility relation C is defined as follows:

Definition 8. For any compatibility relation C and any i-formula ϕ, the compatibility relation C *updated with* $i\!:\!\phi$, denoted by $C \circ i\!:\!\phi$, is defined as:

$$\bigcup_{\mathsf{upd} \in \mathsf{upd}_\phi} \left(\bigcup_{c \in C} i\!:\!\mathsf{upd}(c) \right)$$

Notice, again, that $C \circ i\!:\!\phi$, is defined only when $i\!:\!\mathsf{upd}(c)$ is defined for all $c \in C$. When $C \circ i\!:\!\phi$ is defined, the update operation satisfies the seven KM postulates [10]. For lack of space we omit the details of the comparison.

When the knowledge base is split into different contexts, there are other new interesting properties of the update operator concerning the relation between the different contexts. We show that our update operation has the following important property:

> **Locality of effects:** *If the knowledge base i is independent from the knowledge base j, then an update in i does not affect j.*

To prove this property, we need to define when a context j is *independent* from a context i. To this purpose, we introduce the notion of *bridging path*. Given a set of bridge rules BR a *bridging path* from i to j is a sequence of bridge rules $br_1, \ldots br_n$ of BR, such that the consequence of br_n is in j and the indexes of every premise of br_k $(1 \leq k \leq n)$ is equal to i or to the index of the consequence of a br_h with $h < k$. Intuitively a bridging path contains the sequence of bridge rules that one has to apply to derive, in context j, the consequences of the assumptions in the context i. We say that j is independent from i if there is no bridging path from i to j.

Proposition 2 (Locality of effects). *If BR does not contain a bridging path from i to j, then $C \models i\!:\!\phi$ if and only if $C \circ j\!:\!\psi \models i\!:\!\phi$.*

4 Computing contextual updates

In this section, we propose a sound and complete procedure, based on Natural Deduction [16], that computes the chains in the compatibility relation $i\!:\!\mathsf{upd}(c)$. This procedure builds a tree of literals in contexts, called *propagation tree*. A propagation tree for $i : \mathsf{upd}(c)$ is build by starting form the main effect of upd in i, by applying bridge rules to determine derived effects, and by making assumptions that certain literals in contexts are not changed. With no loss of generality we define the procedure w.r.t a set of bridge rules BR of the form $i_1\!:\!\phi_1, \ldots i_n\!:\!\phi_n \rightarrow i\!:\!\psi$. where ϕ_k $(1 \leq k \leq n)$ and ϕ are disjunctions of literals.

Definition 9 (Propagation tree). A propagation tree is a finite tree T, such that each node of T is associated either with a set of expressions of the form $i\!:\!l$ and $i\!:\!l^h$, or with FAIL or OK. A branch t of T is *closed* if it contains a node with OK, or with FAIL, otherwise t is *open*. T is *closed* if all its branches are closed, otherwise T is *open*. t is a *valid* branch if it is closed and contains a node labelled with OK. T is *valid* if it contains a *valid* branch.

Expressions of the form $i:l^h$ are called *hypothesis* or *assumptions*. As a point of notation, we use $i:l^{(h)}$ to denote both $i:l$ and $i:l^h$. We also need the concept of *local entailment*. A set of indexed literals S (i.e. a set of formulas of the form $i:l$) *locally entails* another set of indexed literals S', if, for each $i:l \in S'$, either $i:l^{(h)} \in S$ or $i:\bot \in S$.

Definition 10 (Update propagation tree). Let c be a chain satisfying a set BR of bridge rules, and upd a local update on L_i. The *propagation tree* for $i:\mathsf{upd}$ in c is the tree T, built starting from the root node labelled with $i:me(\mathsf{upd})$ and applying the following rules to any open branch t of T in this order:

R0 (Derive False): If $i:l^{(h)} \in t$ and $i:\neg l^{(h)} \in t$ and $i:\bot \notin t$, then extend t with $i:\bot$;

R1 (Close with fail): If $i:\bot \in t$ and t does not contain any assumption, then extend t with FAIL;

R2 (Apply bridge rules): Let $\{br_1, \dots br_n\} \subseteq BR$ be the set of bridge rules such that t locally entail the premises of br_k, and does not locally entails the consequence of br_k. For any $1 \leq k \leq n$ let $i_k : \phi_k$ be the consequence of br_k. Extend t with the nodes labelled with the set $\{i_1:l_1 \dots, i_n:l_n\}$ for any choice of $i_k:l_k \in i_k:\phi_k$, with $1 \leq k \leq n$.

R3 (Discharge assumptions): Let N be a node of t containing a literal $i:l^h$ and such that all the branches containing N contain also $i:\bot$. Let $\{i_1:l_1^h, \dots, i_n:l_n^h\}$ be the set of hypothesis occurring in these branches between N and the first occurrence of $i:\bot$ below N. For each of such nodes N, extend t with the brother nodes $N_1, \dots, N_n$ of N, where N_k is labelled with the following set:

$$\{i:l_1^h, \dots, i:l_k^c, \dots, i:l_n^h\}$$

R4 (Make minimal assumptions): Let $br \in BR$ be a bridge rule, such that

1. t contains at least one of the indices appearing in br;
2. t does not locally entail the consequence of br;
3. there is a minimal set of literals $c \models H$, such that $t \cup H$ locally entails the premises of br;

then, extend t with the leaf node labelled with $\{i:l^h | i:l \in H\}$.

R5 (Close with success): If no more rules can be applied extend t with OK.

Definition 11 (Chains of a tree). Let T be a closed propagation tree for $i:\mathsf{upd}(c)$ and t a valid branch of T. The *chain* corresponding to t, denoted by c^t, is obtained by changing c as follows: for all $i:l \in t$ and for all $m \in c_i$, if l^c in m, then replace l^c with l in m. The set of chains of T is given by the set of chains c_t corresponding to a valid branch t of T, that satisfies the bridge rules BR.

Example 6. Consider once again Example 1: suppose to have the initial situation where both $Mr.1$ and $Mr.2$ see no balls Mr.2 doesn't see $Mr.1$ in front of him. The chain that formalizes this situation is ⟨[|], [| |]⟩. Suppose now that $Mr.1$ updates its beliefs by adding a ball on the left slot of the box. The minimal effects on the beliefs of $Mr.2$ is that he believes that there is a ball in the right slot of the box. Therefore the resulting chain should be ⟨[•|], [|•|]⟩. This unique solution is computed in the propagation tree shown in the side picture. We start with the root node containing the main effect of the update $1\!:\!L$. The first applicable rule is R4 and we therefore make the assumption $2:\neg M$, which activates the bridge rule BR1, obtaining $2\!:\!R$. Again the first applicable rule is R4: possible minimal assumptions are $1\!:\!\neg R$ and $2\!:\!\neg L$. The former activates bridge rules BR7 and the latter BR13 and no more rules can be applied except R5.

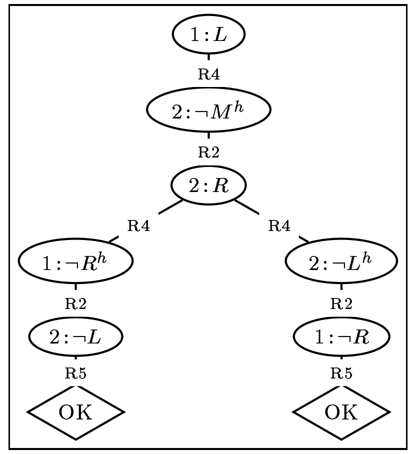

The bridge rules BR1-16 of Example 1, have a very special form. Namely all the consequences of the bridge rules are literals. Furthermore in this example we don't have any local constraint, i.e. some constraint that must be satisfied inside a single context. In the next example we slightly enrich the example in order to show the behavior of the algorithm in more complex cases.

Example 7. Let us modify Example 1, by introducing the local constraint establishing that $Mr.1$ can see at most one ball, formalized by the bridge rule with no premises BR0, shown in the side box. We also allow Mr_1 to affects $Mr.2$ beliefs about the man being in front of him. We therefore rephrase bridge rules BR1–4 with the bridge rules BR1′–BR4′ shown in the side box (the symbol $\supset$ is the implication). To simplify the construction of the tree we drop bridge rules BR5–8 and BR13–16. Consider the initial chain $c =$ ⟨[•|], [•| |M]⟩, formalizing the beliefs of $Mr.1$ and $Mr.2$ in the situation shown in the left box of Figurefig-example. Figure 3 shown the propagation tree computing the two chains in $2:\mathsf{del}(M)(c)$. Namely $c_1 =$ ⟨[|•], [•| |]⟩ and $c_2 =$ ⟨[•|], [|•|]⟩,

(BR0)		$\rightarrow$	$1\!:\!\neg L \vee \neg R$
(BR1')	$1\!:\!L$	$\rightarrow$	$2\!:\!M \supset L$
(BR2')	$1\!:\!L$	$\rightarrow$	$2\!:\!\neg M \supset R$
(BR3')	$1\!:\!R$	$\rightarrow$	$2\!:\!M \supset R$
(BR4')	$1\!:\!R$	$\rightarrow$	$2\!:\!\neg M \supset L$

Theorem 1 (Soundness). *Let T be the update propagation tree for $i\!:\!\mathsf{upd}$ in c. If t is a valid branch of T and $t(c) \models br$, then $t(c) \in i\!:\!\mathsf{upd}(c)$.*

Theorem 2 (Completeness). *For any $c' \in i\!:\!\mathsf{upd}(c)$ there is a valid branch t in the update propagation tree of $i\!:\!\mathsf{upd}$ in c, such that $t(c) = c'$*

The proofs of Theorems 1 and 2 can be found in [5], and an outline will be included in the full paper.

A detailed study of the complexity of the algorithm is out of the scope of this paper, and it is part of the future work. However, we can observe that the dimension of the tree generated by the algorithm depends on the bridge rules and not from the dimension of the

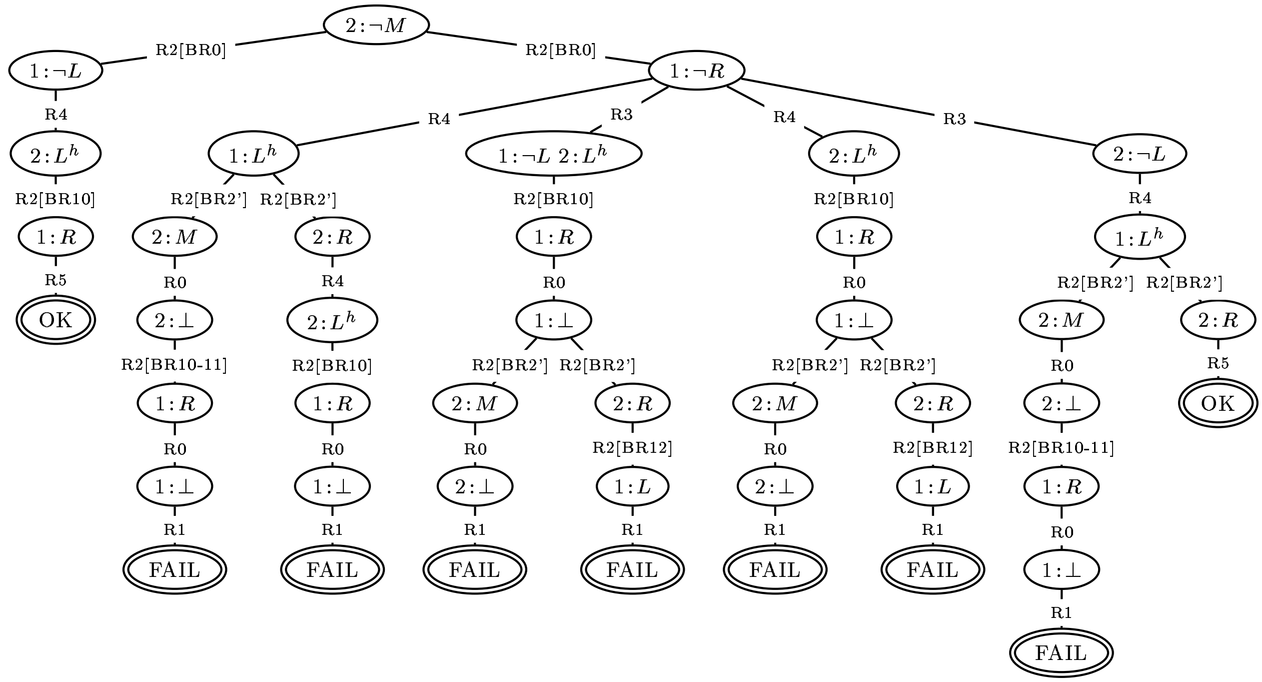

Figure 3: Propagation tree for $2\!:\!\mathsf{del}(M)(c)$

knowledge bases. A first (rough) upper-bound on the ramification rate of the propagation tree is equal to the maximum between m^{n-1} and 2^p (where m is the number of bridge rules, n is the maximal number of disjunctions contained in the consequence of a bridge rule, and p is the different propositional letters appearing in the premises of the bridge rules).

5 Related work

In the area of the formalisms for contextual reasoning, based either on McCarthy or on Giunchiglia's intuitions, this work is, to the best of our knowledge, the first attempt to add some form of "animation" to contexts.

With respect to the well developed area of belief revision and update, we do not introduce new concepts. Our target is complementary as we are more concerned with study of the interactions among local updates in a distributed setting. The approach proposed in this paper is analogous to the MCD (Minimal Change with Maximal Disjunctive Inclusion [21]) based on PMA (Possible Models Approach [20]). The main difference being that we deal with a set of interdependent knowledge bases and not with a unique knowledge base. A second difference is that the knowledge base, in this paper, is represented semantically as a set of propositional models and not as a set of formulas in certain cases this is not possible.

The approaches to the problem of revision of multi-agent beliefs (MABR) described in [11, 19, 12, 13] are related. Kfir-Dahav and Tennenholtz in [11] define a framework for MABR, where agents are required to be consistent on their shared beliefs. In this approach, agents' beliefs are represented by a Knowledge Base on a propositional language L. The primitive propositions of L are partitioned in two sets, *private propositions* and *shared propositions*. Agents must agree on their beliefs on shared propositions. When an agent observes a new fact ϕ, he updates his beliefs accordingly (by any revision strategy); this might involve changes both on private and shared propositions, and, therefore, updates to other agents beliefs. Such behavior is captured by our approach. Intuitively, a proposition p shared by two agents, i and j, can be represented by the bridge rules, $i : p \rightarrow j : p$ and $j : p \rightarrow i : p$. In addition, our approach allows a more fine grained coordination that enables us for instance, to represent asymmetric shared propositions with the bridge rule $i : p \rightarrow j : p$. These propositions represent the fact that changes propagate in the direction "from i to j" and not in the other direction.

The proposal of Liu and Williams in [13] for MABR is based on the paradigm of shared knowledge structure. In this paradigm, the knowledge base of an agent i, KB_i, is partitioned in private knowledge and accessible knowledge. The latter is further partitioned depending on the different degree of accessibility to the other agents. In particular, for each agent j the knowledge base $K_{ass}(i, j) \subseteq KB_i$ is the portion of i's accessible knowledge by agent j. Similarly to [11], $K_{ass}(i, j)$ can be represented in our approach by means of bridge rules, namely the fact $p \in K_{ass}(i, j)$ corresponds to the bridge rule $i : p \rightarrow j : p$.

Van der Meyden in [19] proposes the concept of *Mutual Belief Revision*. Mutual belief revision is the process by which a set of agents revise their beliefs and their mutual beliefs as a consequence of a common observation. At a given time, each agent i is associated with its beliefs, and the beliefs about other agents, (represented by a set of trees isomorphic to a K45 Kripke structure) and a local update function ρ_i. The local update function is supposed to be common knowledge. Roughly, the global update function is defined in terms of iterative composition of local update functions. This work has at least two points in common with our approach. First, update is defined on semantics structures, rather than on set of sentences. Second, global update is the result of the combination of local updates (performed by each agent). The main difference between our approach and the one proposed in [19] is that we have an explicit language to express coordination between agents: bridge rules. Coordination in [19] is instead an underlying assumption which states that local revision functions of each agent are mutually believed. This means that the updates performed by agent i on the beliefs of agents j are the same as that performed by agent j itself. As noticed in [12], this is not a reasonable hypothesis in an heterogeneous distributed system.

6 Conclusions

In this paper we propose a first attempt to define an update operator for contexts. The proposal relies on the formalisms of local model semantics proposed by Giunchiglia and Ghidini in [6]. Technically, we propose an implicit definition of update in context, and provide an algorithm that computes it. We furthermore compare our approach with the current trends in the area of Multi Agent Belief Revision. This work is motivated by providing well founded algorithms for update in distributed and heterogeneous knowledge bases, in the absence of a global representation of the knowledge. Future work will involve discovering an optimized version of the algorithm, studying its complexity, and implementing it in a distributed agent based architecture.

References

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

[2] M. Benerecetti, P. Bouquet, and C. Ghidini. Contextual Reasoning Distilled. *Journal of Theoretical and Experimental Artificial Intelligence*, 12(3):279–305, July 2000.

[3] S. Buvač, V. Buvač, and I. A. Mason. Metamathematics of Contexts. *Fundamentae Informaticae*, 23(3):412–419, 1995.

[4] S. Buvač and I. A. Mason. Propositional logic of context. *Proc. of the 11th National Conference on Artificial Intelligence*, 1993.

[5] A. Donà. Aggiornamento di dati interdipendenti nei multidatabase: modello formale e metodo di calcolo. Master's thesis, Università degli Studi di Verona, Facoltà di Scienze Matematiche, Fisiche e Naturali, Corso di Laurea in Informatica, March 2000. Available as IRST-Technical Report 0002-15, IRST, Trento, Italy.

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

[7] 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.

[8] F. Giunchiglia and L. Serafini. Multilanguage hierarchical logics, or: how we can do without modal logics. *Artificial Intelligence*, 65(1):29–70, 1994. Also IRST-Technical Report 9110-07, IRST, Trento, Italy.

[9] A. Herzig. Logics for belief base updating. In Didier Dubois, Dov Gabbay, Henri Prade, and Philippe Smets, editors, *Handbook of defeasible reasoning and uncertainty management*, volume 3 - Belief Change, pages 189–231. Kluwer Academic Publishers, 1998.

[10] H. Katsuno and A. Mendelzon. On the difference between updating a knowledge base and revising it. In J. F. Allen, R. Fikes, and E. Sandewall, editors, *KR'91: Principles of Knowledge Representation and Reasoning*, pages 387–394. Morgan Kaufmann, San Mateo, California, 1991.

[11] N. E. Kfir-Dahav and M. Tennenholz. Multi-agent belief revision. In Yoav Shoham, editor, *Theoretical Aspects of Rationality and Knowledge: Proceedings of the Sixth Conference (TARK 1996)*, pages 175–196. Morgan Kaufmann, San Francisco, 1996.

[12] W. Liu and M. A. Williams. A framework for multi-agent belief revision, part i: The role of ontology. In *Proceedings 12th Australian Joint Conf. on Artificial Intelligence*, volume 1747 of *Lecture Notes in Computer Science*, pages 168–179, Sydney, Australia, December 1999. Springer.

[13] W. Liu and M. A. Williams. A framework for multi-agent belief revision, part ii: A layered model and shared knowledge structure. In *Linköping Electronic Articles in Computer and Information Science*, volume 5(21), 2000.

[14] B. Malheiro and E. Oliveira. Solving conflicting beliefs with a distributed belief revision approach. In *Proceedings of the Int. Joint Conf. IBERAMIA'2000/SBIA'2000*, pages 146–155, S. Paulo, Brazil, November 2000.

[15] J. McCarthy and S. Buvač. Formalizing Context (Expanded Notes). In A. Aliseda, R.J. van Glabbeek, and D. Westerståhl, editors, *Computing Natural Language*, volume 81 of *CSLI Lecture Notes*, pages 13–50. Center for the Study of Language and Information, Stanford University, 1998.

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

[17] L. Serafini and C. Ghidini. Context Based Semantics for Information Integration. In P. Bonzon, M. Cavalcanti, and R. Nossum, editors, *Formal Aspects of Context*, volume 20 of *Applied Logic Series.* Kluwer Academic Publishers, July 2000. Also IRST-Technical Report 9609-02, IRST, Trento, Italy.

[18] L. Serafini and F. Giunchiglia. ML Systems: A Proof Theory for Contexts. Technical Report 0006-01, ITC-IRST, Trento, Italy, 2000. To appear in the Journal of Logic Language and Information.

[19] R. van der Meyden. Mutual belief revision (preliminary report). In J. Doyle, E. Sandewall, and P. Torasso, editors, *Principles of Knowledge Representation and Reasoning: Proc. Fourth International Conference (KR '94)*, pages 595–606, 1994.

[20] M. Winslett. Reasoning about action using a possible models approach. In Reid Smith and Tom Mitchell, editors, *Proceedings of the Seventh National Conference on Artificial Intelligence*, pages 429–450, Menlo Park, California, 1988. AAAI Press.

[21] Yan Zhang and Norman Y. Foo. Updating knowledge bases with disjunctive information. In *AAAI*, pages 562–568, 1996.