Contextual Reasoning Distilled

M. Benerecetti[1], P. Bouquet[1], and C. Ghidini[2]

[1] Dip. di Informatica e Studi Aziendali,
Università di Trento
Via Inama 5, Trento, Italy
Phone number: +39 0461 88 2135
{bene,bouquet}@cs.unitn.it

[2] Dept. of Computing and Mathematics,
Manchester Metropolitan University
Chester Street, Manchester M1 5GD, U.K.
Phone number: +44 (0)161 247 1556
C.Ghidini@doc.mmu.ac.uk

Abstract

In this paper we provide a foundation of a theory of contextual reasoning from the perspective of a theory of knowledge representation. Starting from the so-called metaphor of the box, we firstly show that the mechanisms of contextual reasoning proposed in the literature can be classified into three general forms (called *localised reasoning*, *push and pop*, and *shifting*). Secondly, we provide a justification of this classification, by showing that each mechanism corresponds to operating on a fundamental dimension along which context dependent representations may vary (namely, *partiality*, *approximation*, and *perspective*). From the previous analysis, we distill two general principles of a logic of contextual reasoning. Finally, we show that these two principles can be adequately formalised in the framework of *MultiContext Systems*. In the last part of the paper, we provide a practical illustration of the ideas discussed in the paper by formalising a simple scenario, called the *Magic Box* problem.

1 Introduction

The notion of context is widely studied in different areas of artificial intelligence (AI). Perhaps the first reference to context in AI can be traced back to R. Weyhrauch and his work on mechanising logical theories in the interactive theorem prover FOL (Weyhrauch 1980). However, it became a popular issue only in the late 1980s, when J. McCarthy proposed to formalise context as a possible solution to 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). In the same years, D. Lenat and R. Guha introduced an explicit mechanism of contexts in the CYC common sense knowledge base. Guha – under McCarthy's supervision – proposed a logic of context in his Ph.D. dissertation. In this work, several and important concepts (such as the formula *Ist(c,p)*, lifting, entering and exiting contexts) were introduced and formalised. F. Giunchiglia was the first to shift the focus explicitly from context to contextual reasoning in his 1993 paper on *Contextual Reasoning* (Giunchiglia 1993). His main motivation was the problem of locality, namely the problem of modelling reasoning which uses only a subset of what reasoners actually know about the world. The proposed framework, called MultiContext Systems (MCS), was then applied to formalise intensional contexts, in particular belief contexts (Giunchiglia, Serafini, Giunchiglia & Frixione 1993, Cimatti & Serafini 1995, Benerecetti, Bouquet & Ghidini 1998)). We refer the reader to (Akman & Surav 1996) for a good discussion of the work on the formalisation of context in AI.

The interest in context is not limited to AI, though. On the contrary, it is discussed and used in various disciplines that are concerned with a theory of representation. In philosophy of language, the notion of pragmatic context has been used to provide a semantics to indexical (demonstrative) languages at least since J. Bar-Hillel's seminal paper on indexical expressions (Bar-Hillel 1954). Almost twenty years later, D. Kaplan published on the *Journal of Philosophical Logic* his well-known formalisation of a logic of demonstratives (Kaplan 1978). A broader philosophical approach to context was proposed and developed by J. Perry in his papers on indexicals and demonstratives, see (Perry 1997). Another approach, based on situation semantics, was pursued by J. Barwise and others (Barwise 1986, Surav & Akman 1995). Recently, R. Thomason has started working on a type theoretic foundation of context (Thomason 1999). In cognitive science, many authors have proposed theories of mental representation where mental contents are thought of as partitioned into multiple contexts (also called spaces (Dinsmore 1991), mental spaces (Fauconnier 1985), etc.). We only need to mention here that the notion of context is very important for other disciplines such as pragmatics, linguistics, formal ontology (see (Bouquet, Serafini, Brezillon, Benerecetti & Castellani 1999) for a recent collection of interdisciplinary papers on context).

Despite this large amount of work, we must admit that we are very far from a generally accepted theory of contextual reasoning. Even if we restrict the focus to theories of representation and language, the definitions of context that can be found in the literature range from '[...] a location – time, place, and possible world – at which a sentence is said' (Lewis 1980) to '[...] a psychological construct, a subset of the hearer's assumptions about the world' (Sperber & Wilson 1986), to '[the] subset of the complete state of an individual that is used for reasoning about a given goal' (Giunchiglia 1993). This makes quite difficult to find an agreement on what the logical structure of reasoning is when context dependent information is involved. Admittedly, many authors investigated special forms of contextual reasoning, but the issue of contextual reasoning in itself has remained in the background. The situation is such that, to the best of our knowledge, no one has succeeded in putting together all this work on context and contextual reasoning in a single theory. The result has been a fragmentation of interests, methodologies, technical tools.

In this paper we aim at providing a foundation of such a theory of contextual reasoning by distilling its basic principles from what has been done in the past. Starting from the so-called metaphor of the box, we illustrate our approach to contextual representation and reasoning (section 2). In section 3 we address the main issue of the paper in four steps: first, we exploit the basic features of the metaphor of the box to show that the mechanisms of contextual reasoning proposed in the literature can be classified into three general forms (called *localised reasoning*, *push and pop*, and *shifting*), each affecting some element of a context dependent representation (section 3.1); second, we show that the classification into three types of mechanism correspond to operations on three fundamental dimensions along which a context dependent representation may vary (section

3.2); third, we use the results of the two previous sections to distill two principles of a logic of contextual reasoning (section 3.3); finally, we show that MCS formalise these two principles (section 3.4). In the last part of the paper, we introduce and formalise the *Magic Box* problem, a simple scenario in which the mechanisms of contextual reasoning can be illustrated and used (section 4).

2 Setting up the context

It is sort of commonplace to say that any representation is context dependent. By this, it is generally meant that the content of a representation 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. Examples are: the location and time when one says ‘It’s raining’; the Sherlock Holmes stories when one asserts ‘Holmes is a detective’; the situation with respect to which we describe the position of two blocks as $on(x, y)$; the qualification that we mean water to drink when one asks ‘Is there water in the refrigerator’ (of course there is water, any food contains water, but we can’t drink it); and so on. There are many reasons why such further information is left implicit. First, it allows much terser representations of common sense facts about the world. In reasoning, it allows agents to disregard a huge amount of potentially available information and concentrate only on what is relevant to solve a problem in a given circumstance. In linguistic communication, it allows a speaker to rely on information that the receiver is supposed to have about the relevant features of the ongoing and possibly past conversations, and in general on common sense knowledge about the world.

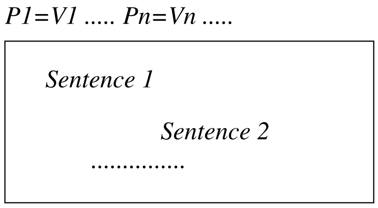

Figure 1: Context as a box

In (Giunchiglia & Bouquet 1997), the notion of context dependence is illustrated by introducing the metaphor of the box (see figure 1). A context dependent representation can be split into three parts: inside the box, a collection of linguistic expressions (be it a single sentence or an entire theory) that describe a state of affairs or a domain; outside the box, a collection of parameters $P_1, \ldots, P_n, \ldots$, and a value V_i for each parameter P_i. The intuition is that the content of what is inside the box is determined (at least partially, and in a sense to be defined) by the values of the parameters associated with that box. For example, in a context in which the speaker is John (i.e. the value of the parameter ‘speaker’ is set to John), any occurrence of ‘I’ will refer to John (we should add a lot of qualifications, but for the moment we will ignore them).

The metaphor can be used to provide a simple illustration of a variety of important (and sometimes controversial) issues in a theory of contextual representation and reasoning. We consider three of them.

The first has to do with the parameters $P_1, \ldots, P_n, \ldots$. What features of context are we to include among $P_1, \ldots, P_n, \ldots$? Is it possible to specify *all* the relevant parameters, or is the collection always incomplete? Is the list of contextual parameters always the same, or is it different from context to context? Theories vary a lot. Kaplan (1978), for example, takes it that all contexts depend on the same collection of parameters, and that this collection is finite (actually, it is a quadruple: a world, a time, a speaker, and a location). Others argue that a representation depends on many other features of context. Lewis, in his 1970 paper on general semantics (Lewis 1970), lists eight parameters; Lenat, in a recent work on the dimensions of context space (Lenat 1999), discusses 12 parameters (he calls them 'dimensions'); many other authors say that the set of contextual dependencies is very large, virtually infinite (Guha 1991, Sperber & Wilson 1986, McCarthy 1993, Giunchiglia & Bouquet 1997), and is different from context to context. The justifications are quite different. On the one hand, Guha and McCarthy (in some of his papers) provide a 'metaphysical' justification, and argue that contexts are *rich objects*, namely objects that can never be completely described. On the other hand, Giunchiglia and Bouquet (1997) propose an 'epistemological' explanation; they argue that *in practice* we can't get a complete list of dependencies because we only have partial knowledge of the world (this facts must be taken into account in an epistemologically adequate theory of representation). In addition, some authors argue that different boxes may have different collections of parameters associated with them. For example, the fact that a block x is on a block y in a situation s can be represented in different ways. In a context in which the situation is left implicit (i.e. it is one of the parameters outside the box), we can use the expression $on(x, y)$ (the value of the parameter tells us to what situation the expressions refers). In a context in which there is no implicit assumption about the situation (i.e. there is no parameter for the situation outside the box), the same fact would be represented with the expression $on(x, y, s)$, where the dependence is made explicit inside the box (Guha 1991, McCarthy 1993, Giunchiglia & Ghidini 1998).

The second issue has to do with the relationship between the parameters $P_1, \ldots, P_n, \ldots$, their value, and the representation inside the box. How do parameters and their value affect the representation of a fact? In what sense a parameter provides implicit information which is to be used in interpreting what is inside a box? Can we get rid of parameters and get a context independent representation of the contents of a box? Various relationships between parameters and representations have been analysed. The relationship is very clear with indexical expressions. Indeed their extension *and* intension is determined by the value of contextual parameters. For example, if the speaker associated with a context is changed from John to Mary, then the content of the pronoun 'I' is modified accordingly. However, there are other possible relationships. Perry, for example, discusses the concept of unarticulated constituent, namely objects that are left implicit in a representation because they can be retrieved from the context. Unlike indexicals, here nothing in the representation indicates that there is a contextual dependency. Consider this example from (Perry 1997). The relation 'raining' is defined between a location and a time. However, if the residents of Z-land never get any information about the weather anywhere else, they can use the sentence 'It's raining' leaving unarticulated the location (i.e. Z-land). In other words, the location is included among the contextual parameters. In AI, the prototypical example of this sort is McCarthy's 'above$-$theory' (McCarthy 1993) (from now, we will refer to this example as the A$-$T example). The notion of unarticulated constituent can be generalised. Indeed, not only arguments of predicates can be left unarticulated. All often, assumptions are left implicit. Some examples: fictional contexts (when we say 'Holmes is a detective' in a context in which we implicitly assume that we are talking about the Sherlock Holmes stories), counterfactual contexts (when one says ‘I would speak a perfect English’ in a context in which he/she implicitly assumes the counterfactual hypothesis ‘If I were born in the UK’). Moreover, contextual parameters may restrict quantification, for example when one says ‘All dogs are sleeping’ in a context in which it is clear that he/she is referring to a particular set of dogs (e.g. the dogs in a room, his/her dogs, and so on).

The third issue has to do with the relationship among boxes: what is the relationship between the parameters of different boxes? How does this relationship affect the relationship between the contents of different boxes? In some cases, the relationship between the parameters of different boxes is very intuitive. For example, if one of the parameters is time (e.g. the day), then the relation is the obvious one (e.g. January 1st, 2000 precedes January 2nd, 2000). For location is slightly more complicated, but still intuitive. However, for other parameters (e.g. beliefs, counterfactual hypothesis, …) the relationship (if any) is less obvious. This has important implications for the issue of contextual reasoning. For example, when the relationship between parameters is well understood, there is a natural way for exporting facts from one context to another. Kaplan’s logic of demonstratives is perhaps the best example of this kind: if ‘Yesterday it was raining’ is true in a context in which time is set to January 2nd, 2000, then ‘Today is raining’ must be true in a context in which time is set to January 1st, 2000. If we had suitable inference rules, ‘Today is raining’ could be derived in the second context starting from the premiss ‘Yesterday it was raining’ in the first only because there is an order between the parameter’s values (from now, we will refer to this example as the Y$-$T example). In other cases, the relation is not obvious, and sometimes there is no relationship at all.

The aim of discussing the three issues above is to set up the context of our analysis of contextual reasoning by clarifying what issues are dealt with in the paper, what are not, and why. Our attitude is the following. On the one hand, we will disregard those aspects that do not affect a theory of contextual reasoning. For example, we do not address the issue of what is the ‘right’ collection of contextual parameters (if any), since the reasoning mechanisms are the same, whatever the collection is. On the other hand, we do not want to make any unnecessary commitment on the issues we deal with. Since our goal is not to propose a particular theory of contextual reasoning, but to provide a general foundation, we try – when possible – to account for the more general theories, and see the others as special cases. Thus, we will allow for infinite collections of parameters, and treat a finite number of parameters as a special case. We will allow for contexts with different collections of parameters, and treat theories in which every context has the same parameters as a special case.

3 Contextual reasoning

The *problem of contextual reasoning* can be intuitively stated as the problem of understanding the general mechanisms that people use to reason with information such that (i) its representation depend on a collection of contextual parameters, and (ii) is scattered across a multiplicity of different contexts.

Mechanisms for contextual reasoning have been studied in different disciplines, though with different goals. However, we still lack a unifying perspective on what a logic of contextual reasoning should do. As a consequence, it is very difficult to see the relationship between the work on context in different disciplines (not to mention the fact that sometimes this relationship is not clear even within the same discipline!). Indeed, there are good pieces of work on utterance contexts, belief (and other intensional) contexts, problem solving contexts, cognitive contexts, and so on, but it's not clear whether they address different aspects of the same problem, or different problems with the same name. In this section, we try to put some order in this situation by addressing the problem of contextual reasoning from a foundational perspective.

3.1 Forms of contextual reasoning

In the past two decades, a large repertoire of mechanisms for contextual reasoning have been identified and formalised. A very partial list includes: *reflection* and *metareasoning* (Weyhrauch 1980, Giunchiglia 1993), *specialisation*, *entering and exiting context*, *lifting*, *transcending context* (Guha 1991, McCarthy 1993, Buvac & Mason 1993), *local reasoning*, *switch context* (Giunchiglia 1993, Bouquet & Giunchiglia 1995), *parochial reasoning*, *context climbing*, and *context initialisation* (Dinsmore 1991), *changing viewpoint* (Attardi & Simi 1995), *reasoning into regions* (Lansky 1991), *focused reasoning* (Hayes-Roth 1991, Laird, Newell & Rosenbloom 1987). The question we address in this section is: What mechanisms of contextual reasoning do they capture? In answering this question, our goal is not to compare formalisms from a technical point of view. What we are looking for is a way of classifying them according to the general mechanism they implement.

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

Localised reasoning. A first general form of contextual reasoning is based on the intuition that some reasoning processes are local to a single box (context), as if the information in that box was the only information available on a given occasion. The idea is that, once we have fixed a collection of contextual parameters and their values, there are reasoning processes that take into account only what is inside a box, disregarding the rest. The most intuitive case of localised reasoning is when an agent is reasoning about a specific and well recognised domain of discourse (say the Italian cuisine, soccer, the Sherlock Holmes stories, the 1960s). Another typical case is problem solving contexts, namely short living representations which contain only the information that a reasoner deems to be relevant to the solution of a particular problem.

Examples of localised reasoning are McCarthy's *reasoning in a context*, Giunchiglia's *local reasoning*, Dinsmore's *parochial reasoning*. However, though they all share a common intuition, the differences among these approaches are quite significant. Some authors think of localised reasoning simply as a way of partitioning a knowledge base according to some pragmatic rule; the expected advantage is that it makes reasoning simpler by reducing the number of potential premisses at each reasoning step. Other authors, however, argue that this 'divide and conquer' approach is not enough. Because of the existing relationships among different contexts, localised reasoning is not simply equivalent to reasoning in a partition of a knowledge base. One must take into account that some facts can be inferred locally only because other facts can be inferred in other contexts (in the Y−T example, 'Today is raining' in the context of January 1st can be inferred from 'Yesterday it was raining' in the context of January 2nd). The second important difference

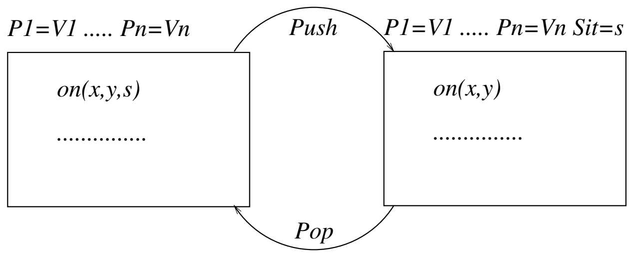

Figure 2: Push and pop

concerns the logic which is used to reason locally. Most authors take it that the logic is the same in every context (reasoning is local, but the logic is global). Others propose that localised reasoning in different contexts may obey different logics (e.g. closed world assumption when checking the train schedule, and deduction when proving theorems in classical logic), see e.g. (Giunchiglia 1993).

Push and 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 (the analogy does not entail that there is an order among parameters, though). 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.

Consider, for instance, the A$-$T example. 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. The idea is that 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 2), whereas pop is the reasoning mechanism which allows us to move back to $on(x, y, s)$ (right-to-left arrow in figure 2). 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

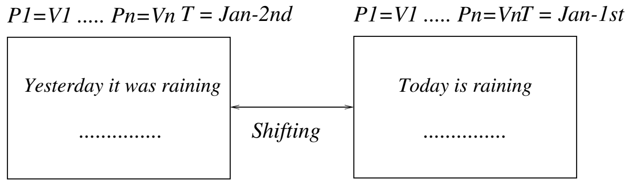

Figure 3: Shifting

and results in the formula $c_0 c : p$, making the dependence on c implicit in the context sequence. This special form of push and pop has been studied by many authors as a separate issue. So, for example, Giunchiglia uses *reflection up* to pop the collection of parameters and *reflection down* to push it; Dinsmore introduces a rule of *context climbing* to pop the collection of parameters, and a rule of *space initialisation* to push it.

Notice that pop is related to the problem of decontextualisation. It can be viewed as a form of (relative) decontextualisation, through which we get rid of parameters and explicitly represent the corresponding information within the box. As such, it is a philosophical minefield. There's been a lot of argument on whether we can keep on popping the collection of parameters until we get rid of *all* contextual dependencies. Philosophers, linguists, psychologists, sociologists, are split into opposite factions. Correspondingly, we find formalisations based on the assumption that a complete decontextualisation is possible (e.g. Dinsmore's space called **base**), and others that deny the existence of such a most general context (e.g. (Guha 1991, McCarthy 1993, Giunchiglia & Bouquet 1997)). Though this issue is very important, in this paper we do not need to commit ourselves to either position. Indeed, from the standpoint of a theory of contextual reasoning, choosing one attitude or the other does not affect the pop mechanism.

Shifting. A third form of reasoning has to do with changing the value of contextual parameters. In other words, unlike push and pop, what changes is not the collection of parameters, but their values. The name 'shifting' is inspired to the concept of shifting in (Lewis 1980). The intuition is that changing the value of contextual parameters shifts the interpretation of what is represented inside the box.

The simplest illustration of shifting is reasoning with indexical expressions. Let us consider in more detail the Y−T example. The fact that on January 1st it is raining can be represented as 'Today is raining' in a context in which time is set to January 1st, while it can be represented as 'Yesterday it was raining' if the parameter is set to January 2nd. As it is shown in figure 3, *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.

Shifting is not limited to indexical expressions. 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 categorisation. 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 Mc-

Carthy formalise a form of shifting using the notion of *lifting*; Dinsmore introduces the notion of *secondary context*; Giunchiglia uses *bridge rules* (though bridge rules, like lifting, are used to formalise other forms of contextual reasoning as well). Needless to say, the differences among these formalisations are many and significant. For example, Kaplan allows shifting between any pair of contexts (this is a consequence of the fact that his theory assumes that the language is the same in all contexts and assigns to each context the same collection of parameters, only with different values), whereas others restrict shifting to pairs of contexts whose parameters are related to each others. A further disagreement is about the interpretation of shifting. Some authors (e.g. Guha) argue that shifting provides a very strong mapping (i.e. logical equivalence) between the contents of different contexts, whereas others (e.g. Giunchiglia) think that this mapping can only be a weak relation of compatibility between formulae of distinct languages whose objective relationship is – at least in practice – out of reach.

3.2 Dimensions of context dependence

The three forms of contextual reasoning we described in the previous section may appear as the result of taking too seriously the metaphor of the box and its basic elements: the representation, the parameters, and their values. Indeed, localised reasoning allows for reasoning within a given representation (i.e. with a fixed collection of parameters and values); push and pop allows for adding or removing parameters from the collection of contextual dependencies; and shifting allows for varying the values of a given collection of parameters.

The goal of this section is to show that the three forms of contextual reasoning we have isolated actually correspond to operating on three fundamental dimensions along which a context dependent representation may vary: *partiality*, namely with the portion of the world which is taken into account; *approximation*, namely with the level of detail at which a portion of the world is represented; and *perspective*, namely with the point of view from which the world is observed. If we succeed in arguing this correspondence between the three forms of contextual reasoning and these three dimensions of representation, then we are in the position of describing a logic of contextual reasoning as the logic of the (formal) relations between partial, approximate, and perspectival representations of the world.

Partiality We say that *a representation is partial when it describes only a subset of a more comprehensive state of affairs*. The intuition is illustrated in figure 4. The circle below represents a state of affairs. From a metaphysical perspective, it may be the world; cognitively, we can imagine that it is the totality of what an agent can talk about. The circles above stay for partial representations of the world, namely representations of portions of it. The figure suggests that there may be some relationships between partial representations (such as overlapping and inclusion), but we do not discuss this aspect here. A sentence such as 'It's raining' taken in isolation, a set of axioms describing the blocks world, a cookbook, a handbook of biology, the Sherlock Holmes stories, are all examples of partial representations.

Perhaps the best example of partial theories in the literature are *micro-theories* (MT) in CYC (Lenat & Guha 1990), each of which represents a small knowledge base about a particular domain (this idea has been refined in (Lenat 1999), where it is proposed that each context is a point in

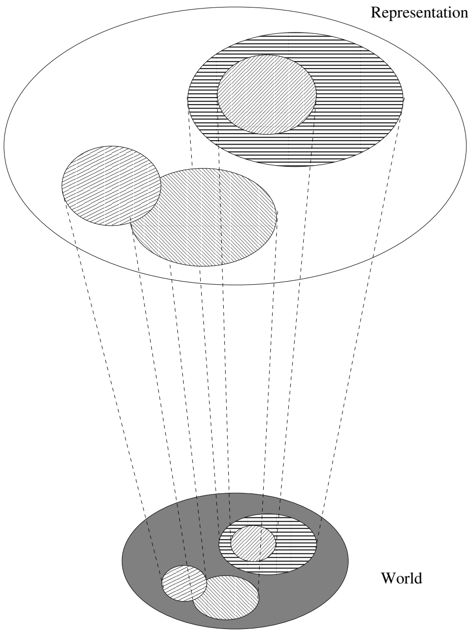

Figure 4: Partiality

a twelve-dimensional space of parameters). Dinsmore's partitioned representations (Dinsmore 1991), and *situations*, as defined in (Barwise & Perry 1983), are other examples.

A different usage of partial theories is in problem solving. In general, given a problem, people seem to be capable of circumscribing what knowlege is relevant to solve it, and disregard the rest. In this case, assumptions on what is relevant act as contextual parameters.

Finally, partial theories are 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 & Wilson 1986), '[a] context in this sense is not limited to the information about the immediate physical environment or the immediately preceding utterances: expectations about the future, scientific hypotheses or religious beliefs, anecdotal memories, general cultural assumptions, beliefs about the mental state of the speaker, may all play a role in interpretation'. However complex, such an interpretation 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.

Approximation We say that *a representation is approximate when it abstracts away some aspects of a given state of affairs*. A description of an office in terms of walls, windows, doors, chairs, tables, plugs, and so on is an approximate representation, because it abstracts away aspects such as the chemical components of furniture or sub-atomic particles. A representation of the blocks world in terms of the binary predicates $on(x, y)$ and $above(x, y)$ is approximate, because aspects such as the situation, the colour of the blocks, their weight, and so on, are abstracted away.

Figure 5 illustrates this idea. The bottom circle represents the world as before. The circles above correspond to possible representations of the world at different levels of approximation. The figure depicts the hierarchy of representations as if each of them covered the entire world.

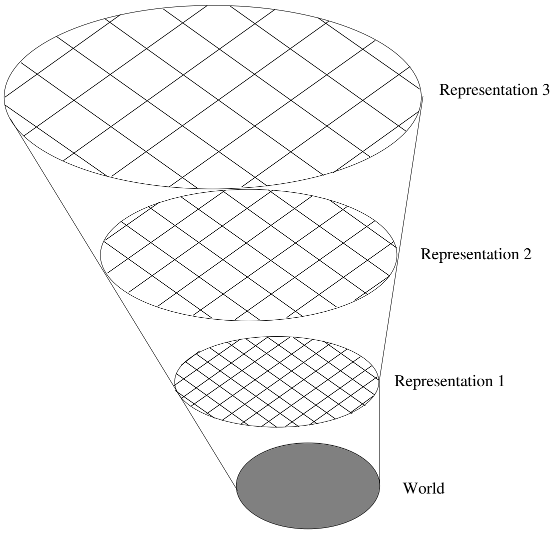

Figure 5: Approximation

However, it should be clear that an approximate representation can also be partial (in the sense we defined above), and therefore it may represent a portion of the world.

This notion of approximation is relative: 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 $above(x, y, s)$, as the first abstracts away the dependence on the situation. Of course, there is the open point of whether there is such a thing as a non approximate representation of a state of affairs. It would be a sort of least approximate representation, namely a representation which is less approximate than anyone else. Even though this issue is of the greatest importance in some areas of philosophy (for example, in a debate on reductionism), we can avoid committing to one position or the other, as we are only interested in the reasoning mechanisms that allow us to switch from a more to a less approximate representation, and not in the epistemological status of the different representations.

Perspective A third dimension along which a representation may vary is perspective. In general, we say that *a representation is perspectival when it encodes a spatio-temporal, logical, and cognitive point of view on a state of affairs*. Figure 6 is a graphical illustration of the idea. In what follows we discuss some intuitive examples.

The paradigmatic case of spatio-temporal perspective is given by indexical languages. Consider purely indexical expressions, such as 'here' and 'now'. A sentences such as 'It's raining (here)(now)' is a perspectival representation because it encodes a spatial pespective (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 even non indexical sentences, such as 'Ice floats on water', encode a perspective, namely a logical perspective. Indeed, 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). That's why Kaplan, for example, includes a world among the features that define a context, and uses this world to interpret the propositional operator

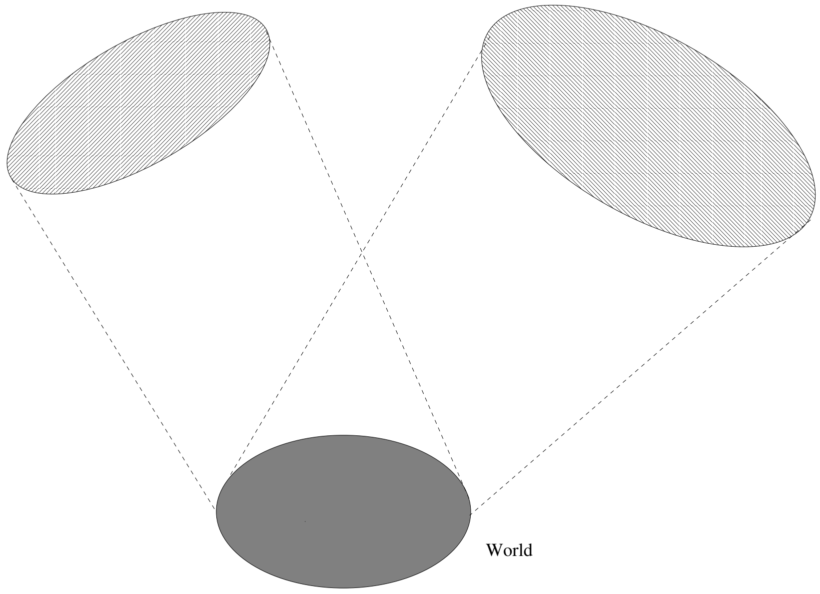

Figure 6: Perspective

'actually'.

Indexicals are not the only expressions that encode a perspective. Suppose, for example, that two agents look at the same object (for example the magic box of figure 7). Because of their different viewpoints, the representation of what they see is completely different, and we can even imagine that the agent named **Side**, who sees a two-sector box, has no idea that the other agent sees a three-sector box. Moreover, 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 representations encode a point of view which includes a collection of beliefs, intentions, goals, and so on. For example, a supporter of team A will refer to supporters of his team as 'friends' and to supporters of team B as 'opponents', whereas B would refer to them the other way around. It goes without saying that a point of view doesn't need to be an individual's point of view. Teams, professional groups, interest groups, societies, cultures, all provide their members with a perspective on their environment. A good example is the way different professional groups within the same organisation represent their knowledge about a domain, making quite difficult the designing of knowledge management systems.

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 three dimensions of partiality, approximation, and perspective:

- localised reasoning is the reasoning mechanism that allows us to exploit partial representations of the world in order to make reasoning more efficient (and, in real world scenarios, to make it possible in practice). Localised reasoning is therefore reasoning that happens when a particular collection of parameters and their values are fixed. It's reasoning in a box, *as if* the box contained all that is needed;

- push and pop is the reasoning mechanism that allows us to vary the degree of approximation by regulating the interplay between parameters outside and representation inside a box. In other words, push and pop is a way of moving from one context to a more (less) approximate one by operating on the collection of parameters, thus making implicit/explicit in the representation some contextual dependencies;

- 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.

3.3 Distilling the principles of a logic for contextual reasoning

The correspondence between forms of contextual reasoning and dimensions of context dependent representations allows us to say that *a logic of contextual reasoning is a logic of the relationships between partial, approximate, and perspectival representations*. The requirements of a general logic of contextual reasoning can thus be stated as follows:

- on the one hand, it must allow for a multiplicity of partial, approximate, and perspectival representations of the world;

- on the other hand, it must formalise the reasoning mechanisms that operate on such representations, namely:
  - localised reasoning, which allows for reasoning within partial representations;
  - push and pop, which allows for reasoning when the degree of approximation at which the world is represented is varied;
  - shifting, which allows for reasoning when the perspective from which the world is represented is changed.

In the past, various logics have been proposed which formalise one aspect or the other of such a logic of contextual reasoning. As we said before, our aim is not to propose a particular logic of contextual reasoning, but to distill the general principles of such a logic. Now the challenge is to find the principles that account for all the requirements we stated at the beginning of the section in their most general form.

Following a traditional view in symbolic AI, we assume that knowledge about a domain can be represented as a logical theory presented as an axiomatic system $\langle L, \Omega, \Delta \rangle$, where L is a formal language (the representation language of the theory), Ω is a set of well-formed formulae of L (the axioms) and Δ is a set of inference rules of L (the inference machinery), *e.g.*, the set of natural deduction rules for L. Reasoning is formalised as inference within the theory (nothing prevents us from assuming that, in general, different contexts may have different inference rules). Now the question we start with is whether it is appropriate to formalise a context as a theory.

Intuitively, different theories may represent different portions of the world (partiality), at different level of detail (approximation), from different perspectives. Moreover, inference within a theory seems to capture very well the idea of localised reasoning. However, this is not enough. Thinking of a context as a theory would not allow us to capture the relationships between different partial, approximate, and perspectival representations. This leads to the following idea: *a context is a theory which is 'plugged' into a structure of relationships with other theories*. In other words, a theory can be part of a context, but we must take into account the fact that there is a structure which acts as a source of additional constraints on what can be derived (what is true) in a context. Intuitively, these constraints are induced by the relationships among the parameters associated with the contexts, and their values. In the Y$-$T example, the structure provides the constraint to the effect that 'Today is raining' cannot be true (derivable) in the context where time is set to January 1st without 'Yesterday it was raining' being true (derivable) in a context where time is set to January 2nd, 2000. In the A$-$T example, the structure provides the constraint to the effect that $on(x, y)$ cannot be true (derivable) in $c(s)$ without $on(x, y, s)$ being true (derivable) in c (i.e. a context in which situations are left implicit).

These ideas can be given a precise model-theoretic and proof-theoretic formulation. Let us introduce some terminology and notation. Suppose $T_1, \ldots, T_n$ is a collection of theories. L_i is the language of the i-th theory, Ω_i is its set of axioms, Δ_i is its inference engine, and *Th(*Ω_i*)* is the transitive closure of T_i. T_i is the theory associated with the context c_i. Let M_i denote the set of *all* possible models of the language L_i. Then $M_{T_i} \subseteq M_i$ is the set of models that satisfy *Th(*Ω_i*)*.

The *first principle of contextual reasoning* (PCR1) can be stated as follows:

First Model-theoretic PCR (PCR1$_{MT}$**)** The set of models that satisfy a context c_i is a subset of M_i.

First Proof-theoretic PCR (PCR1$_{PT}$**)** The set of formulae that can be derived in a context c_i is a subset of L_i.

Indeed, the form of this principle is very general, as it has to account also for non monotonic contexts. In this case, putting additional constraints on a theory may result in changing the set of models that satisfy it, and accordingly the set of formulae contained in the transitive closure. However, when the theory associated with each context is monotonic, the principle PCR1 can be given the following stronger form:

(PCR1'$_{MT}$**):** the set of models that satisfy a context c_i is a subset of M_{T_i};

(PCR1'$_{PT}$**):** the set of formulae that can be derived in a context c_i is a superset of *Th(*Ω_i*)*.

The import of PCR1 may be easily overlooked. It says that a context is a partial, approximate, and perspectival representation in its own right. The language associated with a context c_i is a constraint on what can be expressed in it, and a model of c_i is an interpretation of the language L_i. This means that the semantics of distinct contexts is defined in terms of distinct semantic structures.

PCR1 does not say anything about the relationship between contexts. Let us introduce some further terminology and notation. A *structural constraint* between a pair of contexts c_i and c_j is a relation between the truth of a formula $\phi \in L_i$ and the truth of a formula $\psi \in L_j$. Intuitively, structural constraints are induced by the relationships existing among the parameters (and their values) of c_i and c_j. For instance, the structural constraint of the $\mathrm{Y}-\mathrm{T}$ example is the following: whenever 'Today is raining' (more in general, 'Today$[p]$') is true in a context where time is January 1st, 'Yesterday it was raining' ('Yesterday$[p]$) must be true in a context where time is January 2nd. This constraint is induced by the order of the values for the temporal parameter.

Let R be a structural constraint between ϕ in c_i and ψ in c_j. We say that a model m_i of L_i is *compatible* with a model m_j of L_j with respect to R if, whenever the sentence ϕ is satisfied by m_i, the sentence ψ is satisfied by m_j. For instance, if c_i and c_j are the contexts in which the time is set to January 1st and 2nd, respectively, m_i is *compatible* with m_j if, whenever m_i satisfies a sentence of the form 'Today$[p]$', m_j satisfies 'Yesterday$[p]$'.

Finally, we say that a model m_j of L_j *satisfies a structural constraint* R with the models of L_i if there exists a model m_i of L_i which is compatible with m_j (with respect to R).

The proof-theoretic counterpart is the notion of *structural derivation*. Intuitively, a structural derivation is a derivation of a formula in a context which exploits the structural constraints with other contexts. In the $\mathrm{Y}-\mathrm{T}$ example, 'Today is raining' is structurally derivable in c_i whenever 'Yesterday it was raining' is derivable in c_j.

What we said can be stated as a *second principle of contextual reasoning* (PCR2):

Second Model-theoretic PCR (PCR2$_{MT}$). Only models that satisfy *all* the structural constraints with models of other contexts can be said to satisfy a context.

Second Proof-theoretic PCR (PCR2$_{PT}$). Only formulae that belong to the transitive closure of the union of the theory associated with a context and the formulae derived by structural derivations belong to the transitive closure of the context.

If PCR1 constrains what can be said, PCR2 constrains what is true (what can be derived) in a context. The relation between the two principles should be clear. PCR2 says that there are facts which are true (derivable) in a context because of a relationship existing with some other context; PCR1 says that this relationship is such that nothing which is not locally expressible in a context can be true (derivable) in it.

Structural constraints (derivations) are the model-(proof-)theoretic counterpart of the mechanisms of contextual reasoning described above. In the $\mathrm{Y}-\mathrm{T}$ example, the structural constraint captures a form of shifting; in the $\mathrm{A}-\mathrm{T}$ example, the structural constraints captures push and pop.

Let us briefly discuss if and how PCR1 and PCR2 are formalised in two existing frameworks. Kaplan's logic of demonstratives is such that the set of facts that can be expressed is the same in every context (because the language is the same), and therefore the first principle is trivially satisfied; on the other side, the character of an expression is precisely a way of characterising structural constraints that correspond to a form of shifting (partiality and approximation are not dealt with in his logic). In Buvač and Mason's logic of context, PCR1 is formalised by assuming that the global language is only partially interpreted in each context (this leads them to introduce the notion of a context vocabulary, namely the subset of the language which is interpreted for each context). It is not clear whether their logic fully complies to PCR1, as the formula *ist(*κ*,*ϕ*)* is well defined for any well-formed formula ϕ belonging to this global language, and therefore we can sensibly ask whether the formula *ist(*κ*,*ϕ*)* is true (false, undefined) even for formulae ϕ which are not part of the vocabulary of the context κ. A form of push and pop is 'hardwired' in their logic through the mechanisms of entering and exiting contexts.

3.4 A formalisation of contextual reasoning

To the best of our knowledge, *Local Model Semantics* (LMS) and *MultiContexts Systems* (MCS) are the framework that satisfies these two principles in the most general form. In the remainder of this section we shortly present LMS and MCS and show that they formalise PCR1 and PCR2 in the monotonic case, that is when adding information (axioms and/or constraints) to contexts does not reduce the consequences one can draw. For a more complete presentation of the formalism, the interested reader may refer to (Giunchiglia & Ghidini 1998) and (Giunchiglia 1993). By abuse of notation, we will use the symbol c_i (possibly with different subscripts), to mean either the theory associated with context c_i or a context embedded in a structure of relationships with other contexts.

In LMS, one starts with a family of languages $\{L_i\}_{i\in I}$ (hereafter $\{L_i\}$). Intuitively, L_i is the representation language of a context (or theory) c_i. Each language L_i has its set of models M_i. Every subset M_{T_i} of M_i satisfies a set of formulae, each corresponding to a different choice of the theory T_i associated to c_i. Once the theory T_i associated with c_i is fixed, a model belonging M_{T_i} is called a *local model* of c_i. The local models of c_i are the models of L_i that satisfy the transitive closure of the theory associated with c_i.

Let us consider, in the following, the simple case of two contexts c_1 and c_2 with a structural constraint relating the truth of a formula A_1 in c_1 to the truth of the formula A_2 in c_2 (the case of multiple contexts and multiple structural constraints is a straightforward generalization). Model-theoretically, a structural constraint is represented as a relation (called *compatibility relation*) among sets of local models $\mathbf{c}_1$ and $\mathbf{c}_2$ of the two contexts c_1 and c_2, namely:

$$\mathbf{C} = \{\langle \mathbf{c}_1, \mathbf{c}_2\rangle \mid \text{ if } \mathbf{c}_1 \text{ satisfies } A_1, \text{ then } \mathbf{c}_2 \text{ satisfies } A_2\} \tag{1}$$

Equation (1) states that the sets of local models $\mathbf{c}_1$ and $\mathbf{c}_2$ are compatible if A_2 is true in the set of models $\mathbf{c}_2$ whenever A_1 is true in the set of models $\mathbf{c}_1$ (where the notion of satisfiability of a formula in a (set of) local model is the same as in the theory associated to c_i.) A *model* for a pair of contexts $\{c_1, c_2\}$ is a non-empty compatibility relation $\mathbf{C}$ defined over sets of (local) models of c_1 and c_2.

The notion of *satisfiability* of a formula of a context in a model $\mathbf{C}$ is defined as follows. A formula of the context c_i is satisfied by a model $\mathbf{C}$ if all the local models of c_i $(i = 1, 2)$ belonging to $\mathbf{C}$ satisfy it. To define it formally, we first extend local satisfiability to sets of local models as follows. Given a set of local models $\mathbf{c}_\mathrm{i}$, $\mathbf{c}_\mathrm{i} \models \phi$ if and only if, for all $m \in \mathbf{c}_i$, $m \models_{lc} \phi$, where $\models_{lc}$ is the local satisfiability relation of the theory associated with c_i. Let now $\mathbf{C}$ be a model for $\{c_1, c_2\}$ and ϕ a formula of L_i $(i = 1, 2)$. Then, $\mathbf{C}$ *satisfies* ϕ in c_i, in symbols $\mathbf{C} \models c_i : \phi$, if and only if for all $\langle \mathbf{c}_1, \mathbf{c}_2\rangle \in \mathbf{C},\ \mathbf{c}_\mathrm{i} \models \phi$. *Validity* of a formula ϕ in a context c_i is then defined as expected:

$$\models c_i : \phi \text{ if and only if for all models } \mathbf{C},\ \mathbf{C} \models c_i : \phi$$

Therefore the set M_{c_i} of local models satisfying context c_i can be defined as the set of local models of c_i allowed by some $\mathbf{C}$. Formally:

$$M_{c_i} = \{m \mid m \in \mathbf{c}_i \text{ with } \langle \mathbf{c}_1, \mathbf{c}_2 \rangle \in \mathbf{C} \text{ for some } \mathbf{C}\}$$

and it is easily seen that $M_{c_i} \models \phi$ if and only if $\models c_i : \phi$.

The definitions given above clearly satisfy both the model-theoretic PCRs. Indeed M_{c_i} can only contain models of L_i which satisfy the theory associated with c_i (as required by PCR1'$_{MT}$), as it is built out of local models of c_i, namely M_{T_i}. Moreover, it also satisfies PCR2$_{MT}$. Indeed, since it must belong to a compatibility relation, each model in M_{c_i} is, by construction, a model that satisfies all the structural constraints with models of the other contexts.

The proof-theoretic counterpart is the following. An MCS is a pair $MC = \langle \{c_i\}, \mathit{BR} \rangle$, where $\{c_i\}$ is a set of axiomatic formal theories (namely triples of the form $c_i = \langle L_i, \Omega_i, \Delta_i \rangle$), and *BR* is a set of bridge rules. *Bridge rules* are rules whose premisses and conclusion belong to different contexts. For instance, the bridge rule corresponding to the structural constraint described above would be the following:

$$\frac{c_1 : A_1}{c_2 : A_2}$$

where $c_1 : A_1$ is the premiss of the rule and $c_2 : A_2$ is the conclusion. Obviously, bridge rules are conceptually different from local rules (*i.e.* rules in Δ_i). The latter can be applied only to formulae of L_i, whereas the former have the premisses and the conclusion that belongs to different contexts. Intuitively, bridge rules allow for the MCS version of structural derivations (see section 4 for some examples).

A deduction in an MCS MC is a tree of local deductions, obtained by applying only rules in Δ_i, concatenated with one or more applications of bridge rules (see (Giunchiglia & Serafini 1994) for a technical treatment). Notationally we write $\Gamma \vdash_{\text{MC}} c_i : A$ to mean that the formula A is derivable in context c_i from Γ in the MCS MC. A formula ϕ of c_i is a theorem of MC if it is derivable from the empty set, notationally $\vdash_{\text{MC}} c_i : \phi$.

As a consequence, a formula ϕ which is theorem, i.e. belongs to the transitive closure, of a context c_i can be proved by combining the application of local inference rules of c_i with inferences obtained as consequences of bridge rules. It follows that each theorem of the theory associated with c_i is also a theorem of c_i, but additional theorems may be proved as a consequence of combining applications of bridge rules and local rules. Thus, MCS satisfy PCR1'$_{PT}$. PCR2$_{PT}$ is clearly satisfied as the transitive closure of a context results from a combination of local and structural derivations.

4 The Magic Box problem

In the present section we are concerned with providing an example of contextual reasoning that we can use to briefly illustrate the ideas expressed in the paper. The example is called the Magic Box (MB) problem, and the solution to the problem we propose involves a very simple case of contextual reasoning. Despite its simplicity, we can use the MB problem to show, in a single example, how MCS formalise the three mechanisms of contextual reasoning, and the relationships between these mechanisms and the dimensions of context dependence.

4.1 The scenario

Suppose there are three observers, **Top**, **Side**, and **Front**, each having a partial view of a box as shown in the top part of figure 7. **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.

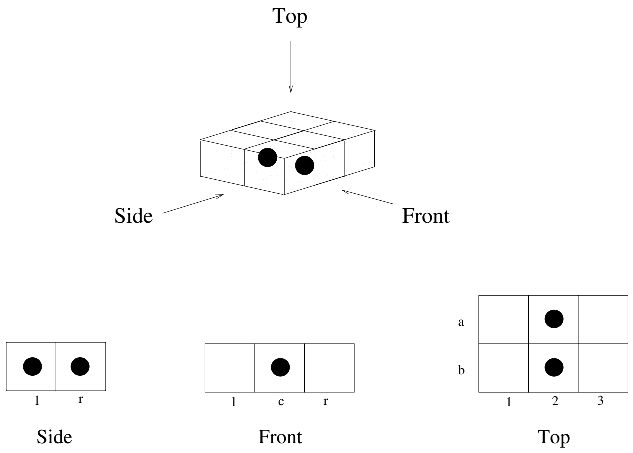

Figure 7: The magic box and its partial views

The box is "magic" and **Side** and **Front** cannot distinguish the depth inside it. The bottom part of figure 7 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. Notice that, in many cases a unique answer of ϵ is not guaranteed as the description of **Side** and **Front**'s partial views is often not enough to determinate **Top**'s view of the box. We will concentrate on the fortunate case depicted in figure 7 in which that is the case.

The computer program ϵ knows that **Top**, **Side**, and **Front** can only see (or talk about) different parts of the box from a specific perspective and what part of the box they can see. Therefore, it also knows how to relate the information coming from the first two observers (**Side** and **Front**) to the representation of the box of the third observer (**Top**) so as to try to build **Top**'s view of the box. Such knowledge is independent from the particular instantiation of the scenario, from the actual position of the balls inside the box and the number of balls in it. Thus, we will keep the knowledge about the relations among the different representations separated from the ground knowledge about the box. We can therefore represents the reasoning process of the computer program in solving the puzzle by means of the four contexts depicted in figure 8. Contexts **Side** and **Front** contain the program's 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; finally, context ϵ contains the knowledge that the computer program has about the game, namely what the relations among the other contexts are. This knowledge represents the fact that the three contexts actually describe the same object: the magic box.

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 in figure 8

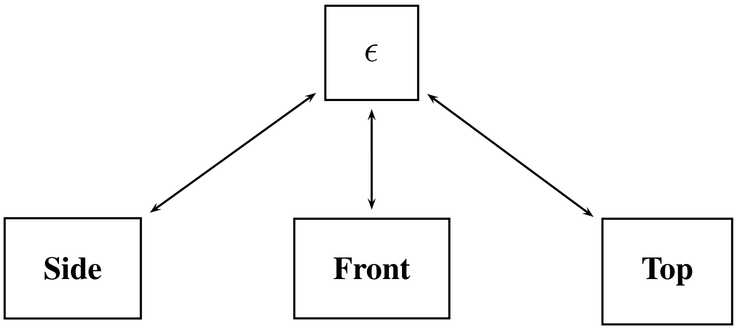

Figure 8: The contexts of the MB scenario

represents different portions of the scenario. For instance context **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 the (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 and pop, and it is therefore related to the different levels of approximation of the different contexts. In this case we say that the representations in **Side**, **Front**, and **Top** are more approximate of the one in ϵ. Indeed 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 perspective are different, the same description (*e.g.* 'A ball is in the left sector') may, thus, have a different meaning in different contexts. In order for ϵ to reason about the relation between the different perspectives, it needs a form of shifting.

4.2 A formalisation of the scenario

Following (Cimatti & Serafini 1995), the first step in formalising the MB example is to introduce the class of languages for the four contexts. Each context has a distinct language, reflecting the fact that each context refers to a different piece of the world and that the world is observed from different perspectives.

Side needs only two atomic propositions to express his/her basic knowledge:

$$AP_{\textbf{Side}} = \{l, r\}$$

meaning that **Side** sees a ball in the *left* sector and *right* sector, respectively, from its point of view. Similarly, **Front** needs three atomic propositions:

$$AP_{\textbf{Front}} = \{l, c, r\}$$

meaning that **Front** sees a ball in the *left* sector, *center* sector and *right* sector, respectively, from its point of view. As $AP_{\textbf{Side}}$ and $AP_{\textbf{Front}}$ are distinct, $l \in AP_{\textbf{Side}}$ is distinct from $l \in AP_{\textbf{Front}}$. **Top** can express its basic knowledge by means of six atomic propositions, one for each sector in the box:

$$AP_{\textbf{Top}} = \{a1, a2, a3, b1, b2, b3\}$$

The corresponding languages, $L_{\textbf{Side}}$,$L_{\textbf{Front}}$ and $L_{\textbf{Top}}$ are the propositional languages built from $AP_{\textbf{Side}}$, $AP_{\textbf{Front}}$ and $AP_{\textbf{Top}}$, respectively. Context ϵ contains the knowledge that the three observers actually talk about the same object from different perspectives, knows what is true in each of them and what the relation among their corresponding contexts is. To account for the variation in the approximation level described above, the language 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 $L_{\textbf{Side}}$ or $L_{\textbf{Front}}$ or $L_{\textbf{Top}}$, and a binary predicate $ist(c, \text{“}\phi\text{”})$, whose intuitive meaning is that formula $\phi \in L_{\text{c}}$ is true in context c.

4.2.1 Formalising the Magic Box with MC systems

Languages The MC system representing the scenario in our case study will contain four contexts **Side**, **Front**, **Top**, and ϵ, with languages $L_{\textbf{Side}}$, $L_{\textbf{Front}}$, $L_{\textbf{Top}}$, and L_{ϵ}, respectively.

Axioms The initial knowledge that each context contains would depend on the particular instantiation of the scenario. Figure 9 shows the knowledge contained in the four contexts, assuming the first observer informs the program about his seeing a ball both in the left and in the right sectors (axioms (1) and (2) of **Side**, respectively.), while the second observer sees one ball in the central sector and no ball in either the left or the right sector from its point of view (axioms (2), (1), and (3) of **Front**). Finally, in this particular instantiation, **Top** has no initial knowledge, and this is the context in which the computer program will try to solve the puzzle. Therefore no (non logical) axiom is in it. The top box (labeled ϵ) in figure 9 shows a formalisation of the knowledge that the program has about the game. For instance, axiom (1) says that **Side** can see a ball in the left position ($ist(Side, \text{“}l\text{”})$) if and only if there is at least a ball in **Top**'s view of the box and it is placed in $a1$ or $a2$ or $a3$ ($ist(Top, \text{“}a1 \vee a2 \vee a3\text{”})$).

Local inference rules The computer program needs to perform reasoning inside each context. For the sake of the example we associate the set of inference rules for propositional logic to each context.

Bridge rules Arrows connecting contexts in figure 9 represent the relations among contexts. Intuitively, they are meant to capture the relations between the two different approximation levels of each observers' context and context ϵ. They state the correspondence between a formula ϕ in each observers' context c and the formula $ist(c, \text{“}\phi\text{”})$ in context ϵ. Each such relation obviously works in a bidirectional way. In particular, if a formula of the form $ist(c, \text{“}\phi\text{”})$ can be proved in ϵ, then the formula ϕ must be provable (be a theorem) in c, and vice-versa. This relation can be

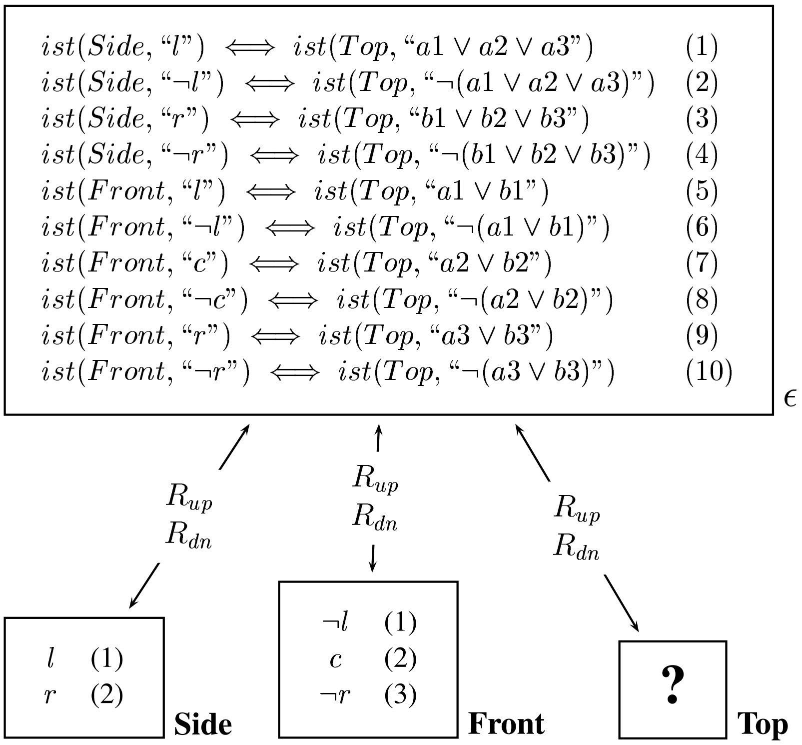

Figure 9: The MC system: local axioms and bridge rules

formally captured by the two bridge rules below:

$$\frac{\mathbf{c} : \phi}{\epsilon : ist(c, \text{“}\phi\text{”})}\ R_{up} \qquad\qquad \frac{\epsilon : ist(c, \text{“}\phi\text{”})}{\mathbf{c} : \phi}\ R_{dn} \tag{2}$$

where **c** can be any of **Side**, **Front** and **Top**. These bridge rules are called *reflection up* (R_{up}) and *reflection down* (R_{dn}), respectively. The formalise the mechanism of push (R_{dn}) and pop (R_{up}).

The solution of the puzzle Given this formalisation, we can show how the contextual reasoning process allows the computer program to solve the puzzle. What we expect it to conclude is that **Top** sees only two balls in the sectors of the central column (see figure 7).

Let us consider one by one the reasoning steps that the computer program can perform. It knows, from the knowledge in context **Side**, that from the left side of the box two balls can be seen (axioms (1) and (2) in **Side**). Intuitively, this means that there must be a ball in at least one sector of the first row a and in at least one sector of the second row b of the complete box. Derivation Π_1 and Π_2 below show a natural deduction style proof of these conclusions. By means of the reflection rule R_{up} between **Side** and ϵ, the reasoner can prove the formulas $ist(Side, \text{“}l\text{”})$ and $ist(Side, \text{“}r\text{”})$ in ϵ. From the information it has about the game, by local reasoning in ϵ (classical modus ponens with axioms (1) and (2) in ϵ), it can conclude $ist(Top, \text{“}a1 \vee a2 \vee a3\text{”})$ and $ist(Top, \text{“}b1 \vee b2 \vee b3\text{”})$. Finally, by applying the reflection rule R_{dn} between **Top** and ϵ, it can conclude $a1 \vee a2 \vee a3$ and $b1 \vee b2 \vee b3$ in **Top**.

$$\Pi_1 = \left\{ \dfrac{\boxed{\dfrac{\dfrac{\boxed{l \qquad (1)}_{\textbf{Side}}}{ist(Side, \text{“}l\text{”})}\, R_{up}}{ist(Top, \text{“}a1 \vee a2 \vee a3\text{”})}\, mp\ (1)}_{\epsilon}}{\boxed{a1 \vee a2 \vee a3}_{\textbf{Top}}}\, R_{dn} \right. \qquad \Pi_2 = \left\{ \dfrac{\boxed{\dfrac{\dfrac{\boxed{r \qquad (2)}_{\textbf{Side}}}{ist(Side, \text{“}r\text{”})}\, R_{up}}{ist(Top, \text{“}b1 \vee b2 \vee b3\text{”})}\, mp\ (3)}_{\epsilon}}{\boxed{b1 \vee b2 \vee b3}_{\textbf{Top}}}\, R_{dn} \right.$$

The information given by **Front**, namely that it sees only one ball in the central sector, should suggest that there cannot be any ball in the first and third column of the box, while there can be a ball in either one of the sectors in the central column. Proofs Π_3, Π_4 and Π_5 show the reasoning steps that can be carried out between **Front** and ϵ so as to map the information coming from the **Front** into the complete description of the box (context **Top**), and conclude $\neg(a1 \vee b1)$, $a2 \vee b2$ and $\neg(a3 \vee b3)$ in **Top**.

$$\Pi_3 = \left\{ \dfrac{\boxed{\dfrac{\dfrac{\boxed{\neg l \qquad (1)}_{\textbf{Front}}}{ist(Front, \text{“}\neg l\text{”})}\, R_{up}}{ist(Top, \text{“}\neg(a1 \vee b1)\text{”})}\, mp\ (6)}_{\epsilon}}{\boxed{\neg(a1 \vee b1)}_{\textbf{Top}}}\, R_{dn} \right. \qquad \Pi_4 = \left\{ \dfrac{\boxed{\dfrac{\dfrac{\boxed{c \qquad (2)}_{\textbf{Front}}}{ist(Front, \text{“}c\text{”})}\, R_{up}}{ist(Top, \text{“}a2 \vee b2\text{”})}\, mp\ (7)}_{\epsilon}}{\boxed{a2 \vee b2}_{\textbf{Top}}}\, R_{dn} \right.$$

$$\Pi_5 = \left\{ \dfrac{\boxed{\dfrac{\dfrac{\boxed{\neg r \qquad (3)}_{\textbf{Front}}}{ist(Front, \text{“}\neg r\text{”})}\, R_{up}}{ist(Top, \text{“}\neg(a3 \vee b3)\text{”})}\, mp\ (10)}_{\epsilon}}{\boxed{\neg(a3 \vee b3)}_{\textbf{Top}}}\, R_{dn} \right.$$

The proof trees above show how the computer program can combine the information contained in the contexts **Side** and **Front** so as to derive information about the possible configurations of the box from the point of view of **Top**. Below is the proof tree that obtains the puzzle solution starting from the conclusions drawn by the proof trees Π_1,...,Π_5. The label $Taut$ on the application of the rule local to context **Top** is essentially a shorthand for a trivial sequence of classical propositional rules. The final step gives the conclusion $ist(Top, \text{“}\neg a1 \wedge a2 \wedge \neg a3 \wedge \neg b1 \wedge b2 \wedge \neg b3\text{”})$ in ϵ, meaning that **Top** sees only two balls in the sectors of the central column (i.e. $a2$ and $b2$).

$$\dfrac{\boxed{\dfrac{\begin{matrix}\Pi_1 \\ \vdots \\ a1 \vee a2 \vee a3\end{matrix} \quad \begin{matrix}\Pi_2 \\ \vdots \\ b1 \vee b2 \vee b3\end{matrix} \quad \begin{matrix}\Pi_3 \\ \vdots \\ \neg(a1 \vee b1)\end{matrix} \quad \begin{matrix}\Pi_4 \\ \vdots \\ a2 \vee b2\end{matrix} \quad \begin{matrix}\Pi_5 \\ \vdots \\ \neg(a3 \vee b3)\end{matrix}}{\neg a1 \wedge a2 \wedge \neg a3 \wedge \neg b1 \wedge b2 \wedge \neg b3}\, Taut}_{\textbf{Top}}}{\boxed{ist(Top, \text{“}\neg a1 \wedge a2 \wedge \neg a3 \wedge \neg b1 \wedge b2 \wedge \neg b3\text{”})}_{\epsilon}}\, R_{up}$$

Remark. In the proof trees above, derivations within the same rectangle are instances of localised reasoning, whereas the proof steps connecting different rectangles are instances of structural derivations. In the example, all the structural derivations are applications of reflection rules.

4.2.2 Modelling the Magic Box with Local Models Semantics

Languages The class of models representing the scenario in our case study is defined over the four languages $L_{\textbf{Side}}$, $L_{\textbf{Front}}$, $L_{\textbf{Top}}$, and L_ϵ associated to the contexts **Side**, **Front**, **Top**, and ϵ.

Local Semantics The *local models* of each context **Side**, **Front**, **Top**, and ϵ are the propositional models of the corresponding language, which satisfy the initial knowledge (axioms) of the context (see figure 9). The *local satisfiability relation* is the standard satisfiability relation $\models$ among propositional languages and propositional formulae.

Notice that we have decided to consider L_ϵ as a propositional language containing 'special' propositional letters $ist(c, \text{“}\phi\text{”})$. This choice enables us to maintain the technical details as simple as possible. Nonetheless, different languages can have different local semantics, and first order semantics might be used for defining the local semantics of L_ϵ. An example of a logic for contextual reasoning involving first order (local) semantics can be found in (Ghidini & Serafini 1998).

Compatibility constraints The model $\mathbf{C}$ for the MB example is a *compatibility relation* containing tuples of the form

$$\langle \mathbf{c}_{\textbf{Side}}, \mathbf{c}_{\textbf{Front}}, \mathbf{c}_{\textbf{Top}}, \mathbf{c}_\epsilon \rangle$$

where each $\mathbf{c}_i$ ($i \in \{\textbf{Side}, \textbf{Front}, \textbf{Top}, \epsilon\}$) is a set of local models for L_i satisfying the following compatibility constraints:

$$\text{if all the } \mathbf{c}_i \in \mathbf{C} \text{ satisfy } \phi \quad \text{then all the} \quad \mathbf{c}_\epsilon \in \mathbf{C} \text{ satisfy } ist(i, \text{“}\phi\text{”}) \tag{3}$$

$$\text{if all the } \mathbf{c}_\epsilon \in \mathbf{C} \text{ satisfy } ist(i, \text{“}\phi\text{”}) \quad \text{then all the} \quad \mathbf{c}_i \in \mathbf{C} \text{ satisfy } \phi \tag{4}$$

Compatibility constraint (3) corresponds to R_{up}. It says that if a formula ϕ is valid in the context labelled by i (is satisfied in all the $\mathbf{c}_i$), then the formula $ist(i, \text{“}\phi\text{”})$ must be valid in the context ϵ. Compatibility constraint (4) corresponds to R_{dn}. It says that if the formula $ist(i, \text{“}\phi\text{”})$ is valid in the context ϵ, then the formula ϕ must be valid in the context labelled by i.

The solution of the puzzle Now we are ready to show how to model the reasoning process of the computer program ϵ in solving the puzzle and obtaining, from the initial knowledge in figure 9, the conclusion $ist(Top, \text{“}\neg a1 \wedge a2 \wedge \neg a3 \wedge \neg b1 \wedge b2 \wedge \neg b3\text{”})$. Remember that from our definition of model for the magic box, all the $\mathbf{c}_{\textbf{Side}}$, $\mathbf{c}_{\textbf{Front}}$, and $\mathbf{c}_\epsilon$ must satisfy the sets of initial axioms in the contexts **Side**, **Front**, and ϵ (see figure 9).

Constraint (3) between contexts **Side** and ϵ and between context **Front** and ϵ tells us that all the $\mathbf{c}_\epsilon$ in a model $\mathbf{C}$ must satisfy also the following formulae:

due to compatibility with **Side**

due to compatibility with **Front**

$$\boxed{\begin{array}{c} ist(Side, \text{“}l\text{”}) \\ ist(Side, \text{“}r\text{”}) \\[4pt] ist(Front, \text{“}\neg l\text{”}) \\ ist(Front, \text{“}c\text{”}) \\ ist(Front, \text{“}\neg r\text{”}) \end{array}}_{\,\epsilon} \tag{5}$$

From the definition of local semantics as propositional semantics, every $\mathbf{c}_\epsilon$ satisfies also the following logical consequences of the initial axioms in figure 9 and the formulae in (5):

$$\boxed{\begin{array}{l} ist(Top, \text{“}a1 \vee a2 \vee a3\text{”}) \\ ist(Top, \text{“}b1 \vee b2 \vee b3\text{”}) \\ ist(Top, \text{“}\neg(a1 \vee b1)\text{”}) \\ ist(Top, \text{“}a2 \vee b2\text{”}) \\ ist(Top, \text{“}\neg(a3 \vee b3)\text{”}) \end{array}}_{\,\epsilon} \tag{6}$$

Constraint (4) between contexts ϵ and **Top** tells us that all the elements $\mathbf{c}_{\mathbf{Top}}$ in a model **C** must satisfy

$$\boxed{\begin{array}{l} a1 \vee a2 \vee a3 \\ b1 \vee b2 \vee b3 \\ \neg(a1 \vee b1) \\ a2 \vee b2 \\ \neg(a3 \vee b3) \end{array}}_{\,\textbf{Side}} \tag{7}$$

Again from the definition of local semantics as propositional semantics, all the $\mathbf{c}_{\mathbf{Top}}$ must satisfy also the following logical consequence of the formulae in (7).

$$\boxed{\neg a1 \wedge a2 \wedge \neg a3 \wedge \neg b1 \wedge b2 \wedge \neg b3}_{\,\textbf{Side}} \tag{8}$$

Finally, constraint (3) between contexts **Top** and ϵ tells us that each $\mathbf{c}_\epsilon$ in **C** must satisfy

$$\boxed{ist(Top, \text{“}\neg a1 \wedge a2 \wedge \neg a3 \wedge \neg b1 \wedge b2 \wedge \neg b3\text{”})}_{\,\epsilon} \tag{9}$$

which means that **Top** sees only two balls in the sectors of the central column (i.e. $a2$ and $b2$).

Steps (5)–(9) are the model-theoretic counterpart of the proof shown at the end of Section 4.2.1. The application of constraints (3) and (4) corresponds to the application of R_{up} and R_{dn}, respectively. The proof that the MC system defined in Section 4.2.1 is sound and complete with respect to the class of models for the magic box is a straightforward generalisation of the soundness and completeness theorem in (Giunchiglia & Ghidini 1998).

It is worth noting that there is a relation between the perspectives of each pair of observers. Intuitively, it depends on the relation between the value of the parameters that describe their perspective on the box. In our formalisation, we chose to (partially) represent such relations explicitly by means of the axioms in ϵ. However, we could have chosen a different formalisation, in which this relation is encoded as bridge rules. For example, the bridge rule:

$$\frac{\textbf{Side} : l}{\textbf{Top} : a1 \vee a2 \vee a3} \ (1a)$$

would represent the shift of perspective from **Side** to **Top** in case side sees a ball in the left sector. This would be an alternative formalisation of shifting in the MB scenario (see (Giunchiglia & Ghidini 1998) for a formalisation of the MB scenario using this kind of bridge rules).

5 Conclusions

The paper is an attempt of providing a foundation of a theory of contextual reasoning. The main steps of this foundation can be summarised as follows. First, we introduce the so-called metaphor of the box, and show that the mechanisms of contextual reasoning proposed in the literature can be classified according to the element of a contextual representation they affect: the representation itself (localised reasoning), the collection of parameters (push and pop), and the value of parameters (shifting). Second, we argued that each of the three forms of contextual reasoning operates on a fundamental dimension of a context dependent representation: partiality, approximation, and perspective. Consequently, we argued that a logic of contextual reasoning is to be thought of as the logic of the relationships among partial, approximate, and perspectival representations of the world. From this we distilled two principles of a general logic of contextual reasoning, both in model-theoretic and proof-theoretic version. These two principles can be used to evaluate the adequacy of any logic of contextual reasoning which has been proposed.

In a sense, this paper is only a preliminary step of the foundation. Indeed, it opens a whole field of research, both philosophical and logical. Our next step will be studying localised reasoning, push and pop and shifting in the framework of MCS. In particular, we are interested in finding the compatibility relation involved in the three reasoning mechanisms and the corresponding bridge rules. This, we hope, will be part of a new approach to a theory of representation in AI and philosophy, in which context will play a crucial role.

References

Akman, V. & Surav, M. (1996). Steps toward formalizing context, *AI Magazine* **17**(3): 55–72.

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

Bar-Hillel, Y. (1954). Indexical Expressions, *Mind* **63**: 359–379.

Barwise, J. (1986). Conditionals and conditional information, *in* E. Traugott, C. Ferguson & J. Reilly (eds), *On Conditionals*, Cambridge University Press, Cambridge (UK), pp. 21–54.

Barwise, J. & Perry, J. (1983). *Situations and Attitudes*, MIT Press, Cambridge, MA.

Benerecetti, M., Bouquet, P. & Ghidini, C. (1998). Formalizing belief report – the approach and a case study, *in* F. Giunchiglia (ed.), *Artificial Intelligence: Methodology, Systems, and Applications (AIMSA'98)*, Vol. 1480 of *Lecture Notes in Artificial Intelligence*, Springer, pp. 62–75.

Bouquet, P. & Giunchiglia, F. (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.

Bouquet, P., Serafini, L., Brezillon, P., Benerecetti, M. & Castellani, F. (eds) (1999). *Modelling and Using Context – Proceedings of the 2nd International and Interdisciplinary Conference*

(9-11 September 1999, Trento, Italy), Vol. 1688 of *Lecture Notes in Artificial Intelligence*, Springer Verlag - Heidelberg.

Buvac, S. & Mason, I. A. (1993). Propositional logic of context, *in* R. Fikes & W. Lehnert (eds), *Proc. of the 11th National Conference on Artificial Intelligence*, American Association for Artificial Intelligence, AAAI Press, Menlo Park, California, pp. 412–419.

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

Gabbay, D. M. (1996). *Labelled Deductive Systems; principles and applications. Vol 1: Introduction*, Vol. 33 of *Oxford Logic Guides*, Oxford University Press, Oxford.

Gabbay, D. M. & Nossum, R. T. (1997). Structured Contexts with Fibred Semantics, *Proceedings of the 1st International and Interdisciplinary Conference on Modeling and Using Context (CONTEXT-97)*, Rio de Janeiro, Brazil, pp. 46–55.

Ghidini, C. & Serafini, L. (1998). Distributed First Order Logics, *Proceedings of the Second International Workshop on Frontiers of Combining Systems (FroCoS'98)*, Amsterdam, Holland. To appear.

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, Chambery, France, 1993, pp. 39–49. Also IRST-Technical Report 9211-20, IRST, Trento, Italy.

Giunchiglia, F. & Bouquet, P. (1997). Introduction to contextual reasoning. An Artificial Intelligence perspective, *in* B. Kokinov (ed.), *Perspectives on Cognitive Science*, Vol. 3, NBU Press, Sofia, pp. 138–159. Lecture Notes of a course on “Contextual Reasoning” of the European Summer School on Cognitive Science, Sofia, 1996.

Giunchiglia, F. & Ghidini, C. (1998). Local Models Semantics, or Contextual Reasoning = Locality + Compatibility, *Proceedings of the Sixth International Conference on Principles of Knowledge Representation and Reasoning (KR'98)*, Morgan Kaufmann, Trento, pp. 282–289. Short version presented at the AAAI Fall 1997 symposium on context in KR and NL. Also IRST-Technical Report 9701-07, IRST, Trento, Italy.

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

Giunchiglia, F., Serafini, L., Giunchiglia, E. & Frixione, M. (1993). Non-Omniscient Belief as Context-Based Reasoning, *Proc. of the 13th International Joint Conference on Artificial Intelligence*, Chambery, France, pp. 548–554. Also IRST-Technical Report 9206-03, IRST, Trento, Italy.

Guha, R. (1991). Contexts: a Formalization and some Applications, *Technical Report ACT-CYC-423-91*, MCC, Austin, Texas.

Hayes-Roth, B. (1991). Making intelligent systems adaptive, *in* K. VanLehn (ed.), *Architectures for Intelligence*, Lawrence Erlbaum Associates, Hillsdale, NJ, pp. 301–321.

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

Laird, J. E., Newell, A. & Rosenbloom, P. (1987). Soar: An architecture for general intelligence, *Artificial Intelligence* **33**(3): 1–4 64.

Lansky, A. (1991). Localized search for multiagent planning, *Proc. of the 12th International Joint Conference on Artificial Intelligence*, pp. 252–258.

Lenat, D. (1999). The Dimensions of Context Space, *Technical report*, CYCorp. `http://www.cyc.com/context-space.{rtf,doc,txt}`.

Lenat, D. & Guha, R. (1990). *Building large knowledge-based systems*, Addison-Wesley, Reading (MA).

Lewis, D. (1970). General Semantics, *Synthese* **22**: 18–67. Reprinted in (Lewis 1983).

Lewis, D. (1980). Index, Context, and Content, *in* S. Kranger & S. Ohman (eds), *Philosophy and Grammar*, D. Reidel Publishing Company, pp. 79–100.

Lewis, D. (1983). *Philosophical papers*, Oxford University Press. Two volumes.

McCarthy, J. (1980). Circumscription - A Form of Non-monotonic Reasoning, *Artificial Intelligence* **13**(1,2): 27–39. Also in V. Lifschitz (ed.), *Formalizing common sense: papers by John McCarthy*, Ablex Publ., 1990, pp. 142–157.

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.

McCarthy, J. (1993). Notes on Formalizing Context, *Proc. of the 13th International Joint Conference on Artificial Intelligence*, Chambery, France, pp. 555–560.

Perry, J. (1997). Indexicals and Demonstratives, *in* R. Hale & C. Wright (eds), *Companion to the Philosophy of Language*, Blackwell, Oxford.

Sperber, D. & Wilson, D. (1986). *Relevance: Communication and Cognition*, Basil Blackwell.

Surav, M. & Akman, V. (1995). Modeling context with situations, *in* P. Brezillon & S. Abu-Hakima (eds), *Working Notes of the IJCAI-95 Workshop on "Modelling Context in Knowledge Representation and Reasoning"*, pp. 145–156.

Thomason, R. (1999). Type theoretic foundation for context, part 1: Contexts as complex type-theoretic objects, *in* P. Bouquet, L. Serafini, P. Brezillon, M. Benerecetti & F. Castellani (eds), *Modeling and Using Context – Proceedings of the 2nd International and Interdisciplinary Conference (9-11 September 1999, Trento, Italy)*, Vol. 1688 of *Lecture Notes in Artificial Intelligence*, Springer Verlag - Heidelberg, pp. 351–360.

Weyhrauch, R. (1980). Prolegomena to a Theory of Mechanized Formal Reasoning, *Artificial Intelligence* **13**(1): 133–176.