

On the dimensions of context dependence: partiality, approximation, and perspective

Massimo Benerecetti¹, Paolo Bouquet², and Chiara Ghidini³

¹ Department of Physical Sciences – University of Naples,
Via Cintia, Monte S. Angelo – I-80126 Napoli, Italy

² Department of Computer and Management Sciences – University of Trento
Via Inama, 4 – I-38100 Trento, Italy

³ Department of Computer Science – University of Liverpool
Liverpool L69 7ZF, U.K.

bene@na.infn.it bouquet@cs.unitn.it chiara@csc.liv.ac.uk

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

1 Introduction

In the last twenty years, the notion of context has become more and more central in theories of knowledge representation in Artificial Intelligence (AI). The interest in context is not limited to AI, though. It is discussed and used in various disciplines that are concerned with a theory of representation, such as philosophy, cognitive psychology, pragmatics, linguistics. Despite this large amount of work, we must confess that we are very far from a general and unifying theory of context. Even if we restrict the focus to theories of representation and language, it is very difficult to see the relationship between different works on contextual reasoning. As an example, there are good pieces of work on utterance contexts, belief (and other intensional) contexts, problem solving contexts, and so on, but it is not clear whether they address different aspects of the same problem, or different problems with the same name.

In this paper we propose to re-read the past work on context as the search for a logic of the relationships between partial, approximate, and perspectival theories of the world. The idea is the following. We start from an very abstract

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

2 Contexts as boxes

In general, a representation is called context dependent when its content cannot be established by simply composing the content of its parts. In addition, one has to consider extra information that is left implicit in the representation itself. In [13], this notion of a context dependent representation is illustrated by introducing the so-called metaphor of the box (figure 1). A context dependent representation has three basic elements: a collection of parameters P_1, \dots, P_n, \dots , a value V_i for each parameter P_i , and a collection of linguistic expressions that provide an explicit representation of a state of affairs or a domain. The intuition is that the content of what is inside the box depends (at least partially, and in a sense to be defined) upon the values of the parameters associated with box. For example, in a context in which the speaker is John (i.e. the value of the parameter ‘speaker’ is set to John), the content (the intension, if you prefer) of the pronoun ‘I’ will be John, but this is not the case in a context in which the speaker is Mary.

$$P_1=V_1 \dots P_n=V_n$$




Fig. 1. Contexts as boxes.

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

context independent representation of the contents of a box? What is the relationship between the parameters of different boxes? How does this relationship affect the relationship between the contents of different boxes?

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

3 Forms of contextual reasoning

Mechanisms for contextual reasoning have been studied in different disciplines, though with different goals. A very partial list includes: *reflection* and *metareasoning* [22, 14], *entering and exiting context*, *lifting*, *transcending context* [16, 20, 5], *local reasoning*, *switch context* [12, 4], *parochial reasoning* and *context climbing* [7], *changing viewpoint* [1], *focused reasoning* [17]). As a matter of fact, it is very difficult to see the relationship between these different works. We try to put some order in this situation by addressing the problem of identifying the general patterns of contextual reasoning, namely the general mechanisms that people use to reason with information (i) whose representation depend on a collection of contextual parameters, and (ii) which is scattered across a multiplicity of different contexts.

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

Expand/Contract. A first general form of contextual reasoning (depicted in Figure 2) is based on the intuition that the explicit representation associated with a specific context does not contain all the facts potentially available to a reasoner, but only a subset of them. As a consequence, depending on the circumstances, the subset which is explicitly taken into account can be expanded (typically because some new input from the external environment makes it necessary to consider a larger collection of facts), or contracted (typically because the reasoner realizes that some facts are not relevant on a given occasion). An example of expansion is the Glasgow-London-Moscow (GLM) example [19, 4]: when reasoning about traveling from Glasgow to Moscow via London, we normally do not include in the problem solving context the precondition that one must be dressed to get on a plane; however, if one's clothes are stolen at London airport, being clothed becomes a relevant precondition for the success of the travel plan, and therefore the original problem solving context must be expanded with facts about social conventions and buying clothes.

In general, expansion and contraction are used to adjust a particular representation to a problem or to a given goal. The way problem solving contexts are built in CYC (using the strategy of lift-and-solve [16]), the mechanism of building appropriate mental spaces [8] or partitioned representations [7], and the process




Fig. 2. Expand/Contract.

of selecting the relevant facts to interpret an utterance [21] are typical examples of this pattern of contextual reasoning.

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




Fig. 3. Push/Pop.

Consider, for instance, the well known AboveTheory scenario, introduced in [20]. The fact that block x is on block y in a situation s is represented as $on(x, y, s)$ in a context c with no parameter for situations. This is because in some cases we want to leave implicit the dependence on the situation s (typically, when we don't want to take situations into account in reasoning). This means that the situation can be encoded as a parameter, and the representation can be simplified to $on(x, y)$. Push is the reasoning mechanism which allows us to move from $on(x, y, s)$ to $on(x, y)$ (left-to-right arrow in figure 3), whereas pop is the reasoning mechanism which allows us to move back to $on(x, y, s)$ (right-to-left arrow in figure 3). Hence, push and pop capture the interplay between the collection of parameters outside the box and the representation inside the box.

It is worth noting that the mechanism of entering and exiting context proposed by McCarthy and others can be viewed as an instance of push and pop.

Suppose we start with a sentence such as $c_0c : p$, whose intuitive meaning is that in context c_0 it is true that in context c the proposition p is true. The context sequence c_0c can be viewed as the reification of a collection of parameters. Exiting c pops the context sequence, and the result is the formula $c_0 : ist(c, p)$, where the dependence on c is made explicit in the representation $ist(c, p)$ ($ist(c, p)$ is the main formula of McCarthy’s formalism, asserting that a p is true in context c); conversely, entering c pushes the context sequence and results in the formula $c_0c : p$, making the dependence on c implicit in the context sequence. Other examples of push/pop are: *reflection up* to pop the collection of parameters and *reflection down* to push it in [14]; the rule of *context climbing* to pop the collection of parameters, and the rule of *space initialization* to push it in [7].

Shifting. Shifting changes the value of one or more contextual parameter, without changing the collection of parameters. The name ‘shifting’ is inspired to the concept of shifting in [18]. The intuition is that changing the value of the parameters shifts the interpretation of what is represented inside the box.




Fig. 4. Shifting.

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

In the literature, we can find different instances of shifting. Kaplan’s notion of *character* is the semantical counterpart of this reasoning mechanism with indexical languages; Guha and McCarthy formalize a form of shifting using the notion of *lifting* [16]; Dinsmore introduces the notion of *secondary context*.

4 Dimensions of context dependence

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

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

Perhaps the more intuitive example of partial theories are domain specific theories. For instance, a theory about the Italian cuisine is partial because it does not provide information about Indian or French cuisine, about soccer, about quantum mechanics. A different usage of partial theories is in problem solving. Given a problem, people seem to be capable of circumscribing what knowledge is relevant to solve it, and disregard the rest. In this case, assumptions on what is relevant act as contextual parameters. Partial theories are also used in theories of linguistic communication. When a speaker says something to a hearer, it is assumed that the latter interprets what the speaker said in some context. According to [21], '[a] context is a psychological construct, a subset of the hearer's assumptions about the world'. Such a context includes the set of facts that the hearer takes to be relevant in order to assign the correct interpretation to what the speaker said. In this sense, it is a partial theory.

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

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

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

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

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

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

A subtler form of perspective is what we call cognitive perspective. It has to do with the fact that many representations encode a point of view which includes a collection of beliefs, intentions, goals, and so on. Cognitive perspective is very important in the analysis of what is generally called an *intensional context*, such as a belief context. John and Mary may have dramatically different beliefs about Scottish climate, even if they represent the same universe of discourse (or portion of the world) at the same level of approximation. We don’t see any other way of making sense of this difference than that of accepting the existence of a cognitive perspective, which is part of what determines the context of a representation.

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

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

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

5 A logic of contextual reasoning: MultiContext Systems

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

MultiContext systems (MCS) [14], and their *Local Model Semantics* (LMS) [9], provide a logic for contextual reasoning based on the principles of *locality* and *compatibility*. These principles impose that:

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

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

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

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




Fig. 5. The magic box and its partial views.

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

Viewpoints. A paradigmatic example of reasoning with viewpoint is the Magic Box (MB) example, developed in [3].

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

[3] describes a formalization of the reasoning process of ϵ in solving the puzzle, by mean of the four contexts depicted in figure 6. Contexts **Side** and **Front** contain the program’s representation of **Side**’s and **Front**’s knowledge;




Fig. 6. The four contexts in the MB example.

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

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

Formally, the specific domains of **Side**, **Front**, and **Top** are described by three different propositional languages $L_{\mathbf{Side}}$, $L_{\mathbf{Front}}$ and $L_{\mathbf{Top}}$ built up from the sets $AP_{\mathbf{Side}} = \{l, r\}$, $AP_{\mathbf{Front}} = \{l, c, r\}$, and $AP_{\mathbf{Top}} = \{a1, a2, a3, b1, b2, b3\}$ of propositional constants (where l means that the observer sees a ball in the *left* sector, c means that the observer sees a ball in the *central* sector, and so on) To account for the specific domain of ϵ , and its shift in the approximation level described above, the language L_{ϵ} contains a set $\{\mathbf{Side}, \mathbf{Front}, \mathbf{Top}\}$ 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 $L_{\mathbf{Side}}$ or $L_{\mathbf{Front}}$ or $L_{\mathbf{Top}}$, and

a binary predicate $ist(c, \phi)$, whose intuitive meaning is that formula $\phi \in L_c$ is true in context c (see [20]). In order to solve the puzzle ϵ needs to relate information contained in different contexts associated with different levels of approximations. In particular [3] needs to formalize the relation denoted as arrows connecting contexts in figure 6. This is done by imposing a compatibility relation between the models of each observers' context c and models of the context ϵ . To state the correspondence between a formula ϕ in each observers' context c and the formula $ist(c, \phi)$ (denoting the same fact at a different degree of approximation) in context ϵ , if a formula of the form $ist(c, \phi)$ is a theorem in ϵ , then the formula ϕ must be a theorem in c , and vice-versa. The different perspectival representations are formalized in [3] by the different (initial) axioms satisfied in each observer's context and the relations between them are explicitly stated as axioms in context ϵ .

Belief contexts. LMS and MCS have been applied to formalize different aspects of intentional contexts, and in particular belief contexts (see e.g.,[15,6]). An example is a puzzle described in [2], where LMS and MCS are used to solve the problem of the opaque and transparent reading of belief reports.

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

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

We are not interested here in the solution of the puzzle (the interested reader may refer to [2]), but in analyzing the representations of the different contexts involved in the formalization.

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

ϵ represents the context containing the beliefs that the program ascribes to itself, A is the context containing the beliefs that the program ascribes to Mr. A ,




Fig. 7. The structure of belief contexts

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

The formalization of the different contexts in figure 7 may vary along the three dimensions of contextual dependence. Focusing on partiality, the different contexts are related to different sets of beliefs. For instance, **A** is the context containing the beliefs that the program ascribes to Mr. *A*, whereas **B** is the context containing the beliefs that the program ascribes to Mr. *B*. Focusing on approximation, we notice that each context in the hierarchy must be able to talk about beliefs contained in each one of the contexts above. In order to do this it needs to make explicit some information that was implicit in the observers' contexts. In particular, it needs to make explicit what beliefs are contained in what context. The relations involving different contexts associated with different degrees of approximations are the one denoted as arrows in figure 6 and are similar to the ones described in the MB example. Focusing on perspective, each of the belief contexts expresses knowledge about the world which depends on the cognitive perspective of the agents, from the point of view of the computer program. For instance, Mr. *B* will refer to Mr. *C* as "the president of the local football team", whereas Mr. *A* will refer to Mr. *C* as Mr. *M*.

Integration of different information sources. LMS and MCS have been applied to formalize the integration of information coming from different information sources. [10, 11] contain the formal definitions and motivating examples. Let us focus on a simple example.

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

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

The representations of the different contexts in figure 8 may have different degrees of partiality, as each database is associated to a specific domain. For instance, the sellers might provide different subsets of fruits and therefore the domains of their databases are different. Focusing on approximation, the domain




Fig. 8. Contexts in the mediator example.

of fruits can be represented at different level of details by different sellers. E.g., database 1 may contain prices for red apples and yellow apples, while database 2 and 3 abstracts away the dependence on the color and do not make this distinction. Focusing on perspective, prices of the different sellers might be not homogeneous, depending on their particular viewpoint. E.g., prices of database 1 don't include taxes, while prices of database 2, 3 and the mediator do.

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

6 Conclusions

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

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

References

1. G. Attardi and M. Simi. A formalisation of viewpoints. *Fundamenta Informaticae*, 23(2–4):149–174, 1995.
2. M. Benerecetti, P. Bouquet, and C. Ghidini. Formalizing Belief Reports – The Approach and a Case Study. In *Proc. of the 8th International Conference on Artificial Intelligence, Methodology, Systems, and Applications*, volume 1480 of *LNAI*, Springer, 1998.
3. M. Benerecetti, P. Bouquet, and C. Ghidini. Contextual Reasoning Distilled. *Journal of Theoretical and Experimental Artificial Intelligence*, 12(3):279–305, 2000.
4. P. Bouquet and F. Giunchiglia. Reasoning about Theory Adequacy. A New Solution to the Qualification Problem. *Fundamenta Informaticae*, 23(2–4):247–262, 1995.
5. S. Buvač and I. A. Mason. Propositional logic of context. *Proc. of the 11th National Conference on Artificial Intelligence*, pages 412–419, 1993.
6. A. Cimatti and L. Serafini. Multi-Agent Reasoning with Belief Contexts: the Approach and a Case Study. In *Intelligent Agents: Proceedings of ATAL'94*, number 890 in LNCS. Springer, 1995.
7. J. Dinsmore. *Partitioned Representations*. Kluwer Academic Publishers, 1991.
8. G. Fauconnier. *Mental Spaces: aspects of meaning construction in natural language*. MIT Press, 1985.
9. C. Ghidini and F. Giunchiglia. Local Models Semantics, or Contextual Reasoning = Locality + Compatibility. *Artificial Intelligence*, To Appear.
10. C. Ghidini and L. Serafini. Distributed First Order Logics. In *Frontiers Of Combining Systems 2*, Studies in Logic and Computation. Research Studies Press, 1998.
11. C. Ghidini and L. Serafini. Information Integration for Electronic Commerce. In *Proceedings of the Workshop on Agent Mediated Electronic Trading (AMET'98)*, volume 1571 of *LNAI*. Springer Verlag, 1998.
12. F. Giunchiglia. Contextual reasoning. *Epistemologia, special issue on I Linguaggi e le Macchine*, XVI:345–364, 1993.
13. F. Giunchiglia and P. Bouquet. Introduction to contextual reasoning. An Artificial Intelligence perspective. In *Perspectives on Cognitive Science*. NBU Press, Sofia, 1997.
14. F. Giunchiglia and L. Serafini. Multilanguage hierarchical logics (or: how we can do without modal logics). *Artificial Intelligence*, 65:29–70, 1994.
15. F. Giunchiglia, L. Serafini, E. Giunchiglia, and M. Frixione. Non-Omniscient Belief as Context-Based Reasoning. In *Proc. of the 13th International Joint Conference on Artificial Intelligence*, pages 548–554, Chambery, France, 1993.
16. R.V. Guha. Contexts: a Formalization and some Applications. Technical Report ACT-CYC-423-91, MCC, Austin, Texas, 1991.
17. J. E. Laird, A. Newell, and P.S. Rosenbloom. Soar: An architecture for general intelligence. *Artificial Intelligence*, 33(3):1–4 64, 1987.
18. D. Lewis. Index, Context, and Content. In S. Kranger and S. Ohman, editors, *Philosophy and Grammar*, pages 79–100. D. Reidel Publishing Company, 1980.
19. J. McCarthy. Overcoming an Unexpected Obstacle. Unpublished, 1991.
20. J. McCarthy. Notes on Formalizing Context. In *Proc. of the 13th International Joint Conference on Artificial Intelligence*, 1993.
21. Dan Sperber and Deirdre Wilson. *Relevance. Communication and Cognition*. Basil Blackwell, 1986.
22. R.W. Weyhrauch. Prolegomena to a Theory of Mechanized Formal Reasoning. *Artificial Intelligence*, 13(1):133–176, 1980.