A Context-Based Framework for Mental Representation

Fausto Giunchiglia (fausto@irst.itc.it)
ITC-IRST – 38050 Povo, Trento (Italy) and
Dept. of Computer and Management Science, University of Trento
Via Inama, 5 – 38100 Trento (Italy)

Paolo Bouquet (bouquet@cs.unitn.it)
Dept. of Computer and Management Science, University of Trento
Via Inama, 5 – 38100 Trento (Italy)

Abstract

In this paper we present a context-based family of formal systems (called MultiContext Systems) which we propose to use as a formal framework for a theory of mental representation. We start with an intuitive notion of context as a subset of the complete (cognitive) state of an individual. Then we introduce two general principles which we believe are at the core of any logic of context, namely the principles of locality and compatibility. We show how these principles can be formalized in the framework of MultiContext Systems, and argue that this conceptual/logical framework can be used to account for a variety of phenomena in a theory of mental representation. Finally, we compare our framework with previous work.

Introduction

Context is one of the most interdisciplinary notions in the contemporary scientific debate on cognition. In general, context is viewed as a key notion in the explanation of how our mental contents are (logically) structured. As a consequence, context is part of the explanation of how these mental contents are used in common sense reasoning, natural language understanding, learning, problem solving, and so on. The interest in the notion of context is proved by the growing number of scientific events related to it, *e.g.*, the workshops on context at the 13th and 14th International Joint Conference on Artificial Intelligence (Brezillon, 1993; Brezillon & Abu-Hakima, 1995), the workshop on context at the last European Conference on Cognitive Science (Bouquet et al., 1997), the First Interdisciplinary Conference on Context (Brezillion, 1997), the AAAI-97 Fall Symposium on Context in Natural Language and Knowledge Representation (Buvac & Ivanska, 1997).

Despite the large amount of work and number of events, whether there is a unifying notion of context underlying its usage is still unclear. If we look at some of the most well-known definitions in the literature, we find a wide spectrum of proposals. In philosophy of language, David Lewis defines a context as the "location – time, place and possible world – where a sentence is said" (Lewis, 1980) (a very similar notion of context was proposed by David Kaplan in his well-known paper *On the Logic of Demonstratives* (Kaplan, 1978)). In linguistics, Sperber and Wilson define context as "the set of premises used in interpreting an utterance [...] a psychological construct, a subset of the hearer's assumptions about the world" (Sperber & Wilson, 1986). In mental representation, (Dinsmore, 1991) introduces the notion of context as a key concept in the (intuitive) semantics of partitioned representations. In cognitive psychology, Kokinov defines context as "the set of all entities that influence human (or system's) cognitive behaviour on a particular occasion" (Kokinov, 1995). In AI, since John McCarthy's seminal paper on generality (McCarthy, 1987), a lot of work has been done towards the formalization of context and its logic (Guha, 1991; McCarthy, 1993; Buvac & Mason, 1993); however, as McCarthy himself points out, the goal of this work was not to reach a unique conclusion about what context is, but rather to "introduce contexts as abstract mathematical entities with properties useful in artificial intelligence" (McCarthy, 1993). On top of these different approaches, we also find a lot of work about "belief contexts", "intensional contexts", "social contexts", "reasoning contexts", and so on, where the concept of context is assumed as given.

The goal of this paper is to present a context-based family of formal systems (called MultiContext Systems) and show how it can be used to account for many aspects of a theory of mental representation. The paper goes as follows. First, we introduce the basic principles of our approach, namely the *principle of locality* and the *principle of compatibility*. Second, we briefly discuss how these principles are formalized in MultiContext (MC) System. Then we recall some important aspects of a theory of mental representation and show how they can be easily formalized in our framework. Finally, we critically compare our approach with previous work.

Locality and Compatibility

According to (Giunchiglia, 1993), a context can be defined as "that subset of the complete state of an individual that is used for reasoning about a given goal". The intuition is that reasoning is always "local", *i.e.* involves only a small subset of what an agent's actually knows; this subset is what determines the context of reasoning. This intuition can be mapped into the idea that a context is a *partial* and *approximate* theory of the world which encodes some agent's perspective. It is a partial theory because it does not represent all the information that an agent has about the world, but only a small portion. It is an approximate theory because even this portion is represented only at some level of detail. Some examples of context are: the facts that an agent knows about a specific domain (*e.g.* about money, about another agent's beliefs, about the Italian cuisine, ...); the facts that an agent uses to solve a problem on a particular occasion; the facts that an agent uses to interpret another agent's utterance. A context about money, for example, is not only partial (it is about money and not about the Italian cuisine), but also approximate (using an example from (Guha, 1991), the relationship between an object and its cost can be represented at different levels of detail: as a binary predicate *costs(x,y)* (x costs y) in a naive money theory, as a ternary predicate *costs(x,y,ListPrice)* in a general transaction theory in which we may have different kinds of costs associated with an object; and so on).

In our view, a logic of context is the logic of possible relations between partial and approximate theories. The idea is the following:

- on one side, if we take seriously the idea that a context is a theory (though partial and approximate), then expressiveness (what can be said), denotation, truth and logical consequence must be conceived of as local to a context. This is what we call the *principle of locality*;
- on the other side, the fact that an agent's complete knowledge is represented as a collection of contexts suggests that different kinds of relations must exist among different contexts. This is what we call the *principle of compatibility*.

A first consequence of locality is that different contexts may have different ontologies (intuitively, because each of them represents a different portion of an agent's knowledge). For instance, the ontology of a context describing John's beliefs about the world of Calvin & Hobbs will include children, tigers, cars, schools, friendship, and so on; whereas the ontology of a context describing John's beliefs about mechanics will include different categories of objects. A second consequence is that contexts seem to require a "local" notion of denotation and truth. For instance, 'Hobbes' may refer to a speaking tiger in the context of Calvin & Hobbes adventures, and to a philosopher of the XVIII century in a handbook of modern philosophy. As to truth, John may believe that the sentence "Hobbes is a speaking tiger" is true in the context of Calvin & Hobbes stories and that the sentence "Tigers do not speak" is true in the context of zoology. A third important consequence of the principle of locality is that different logics may be allowed in different contexts. For instance, in a context which represents an agent's reasoning about air travelling, we are likely to require a sort of closed-world assumption (*e.g.* if a flight connection from Glasgow to Moscow does not appear on the flight schedule of Glasgow airport, then the agent will infer that a direct connection does not exist); however, in a context involving phone calls, closed-world reasoning does not seem to apply (if the phone number of the Italian Prime Minister Romano Prodi does not appear on Rome's phone book, we don't want to infer that he does not have a phone number in Rome).

The principle of compatibility captures the intuition that contexts, as partial and approximate theories, can be somewhat related. A few examples. Take a context c_1 to describe John's beliefs on January 31, 1998 and a context c_2 to describe John's beliefs on February 1st, 1998. Then it is reasonable to postulate a relation between c_1 and c_2 such that the truth of the sentence "Today's sunny" in c_1 entails the truth of the sentence "Yesterday was sunny" in c_2. Suppose now that a context c_8 describes John's office on January 31, 1998 and that the sentence *on(computer, table)* is true in it. If (for some reason) we switch to a less approximate description where time is explicitly taken into account, then it's reasonable to expect that it will contain a sentence of the form *on(computer, table, 31-12-98)* (and this as a consequence of the fact that *on(computer, table)* was true in c_8). Imagine now that c_3 describes John's office from his own perspective and that Mary is sitting in front of him on the other side of the table. Suppose c_4 represents (some of) the beliefs that John ascribes to Mary. If c_3 contains the sentence *on(computer, left-side-of(table))* (*i.e.* John believes that there is a computer on the left side of the table), then c_4 is likely to contain the belief *on(computer, right-side-of(table))* (John ascribes to Mary the belief that there is a computer on the right side of the table). Finally, assume that c_7 is the context of a conversation about the furnitures in different offices. In c_7 we could have a sentence like *Ist*(c_8, *on(computer, table)*)[1], whose truth has the fllowing relation with the truh of the sentence *on(computer, table)* in the context c_8 (the context describing John's office): *Ist*(c_8, *on(computer, table)*) is true in c_7 only if *on(computer, table)* is true in c_8.

MultiContext Systems

MultiContext (MC) Systems are a formal framework which formalizes the principles of locality and compatibility. MC Systems were introduced in (Giunchiglia, 1993), and formally refined in other papers, *e.g.*, (Giunchiglia & Serafini, 1994; Giunchiglia & Ghidini, 1998). Here we recall just the main concepts; the interested reader may refer to the cited papers for more details.

In a MC System, a context is formalized as a theory presented as an axiomatic formal system. Let L be a logical language (*e.g.*, a first order language), Ω a subset of L and Δ a set of inference rules defined over L. A context c is a triple $\langle \mathrm{L}, \Omega, \Delta \rangle$, where L is the representation language of c, Ω is the set of facts which are assumed to be initially true in c (also called the axioms of c) and Δ is the set of rules which can be used to infer new facts from Ω (*e.g.*, the rules of natural deduction (Prawitz, 1965)). A formula Φ belonging to the language L of a context c is called a L-formula. Informally, we use the notation $c : \Phi$ to express the statement that the L-formula Φ is true in the context c.

The second element of a MC System is a collection of *bridge rules*, namely rules whose premisses and conclusion belong to different contexts. The general form of a bridge rule is the following:

$$\frac{c_1 : \Phi_1, \ldots, c_n : \Phi_n}{c_{n+1} : \Phi_{n+1}}$$

where $c_1 : \Phi_1, \ldots, c_n : \Phi_n$ are the premisses of the rule and $c_{n+1} : \Phi_{n+1}$ is the conclusion.

Notice that the rules in Δ_i (the inference machinery of a context c_i) can be applied only to L_i-formulae (where L_i is the language of c_i), namely premisses and conclusion always belong to the same context. In a bridge rule, instead, the conclusion always belong to a context which is different from the context(s) of the premisses.

A *MC System* is defined as a pair $\langle Cxt, BR \rangle$, where Cxt is a set of contexts and BR is a set of bridge rules.

[1]The predicate *Ist* was introduced by McCarthy in his paper *Notes on formalizing context* (McCarthy, 1993). The intuition is that *Ist(c,p)* holds in a context when the proposition p is true in the context c. For a formal treatment of the modality *Ist* see (Guha, 1991; Buvac & Mason, 1993).

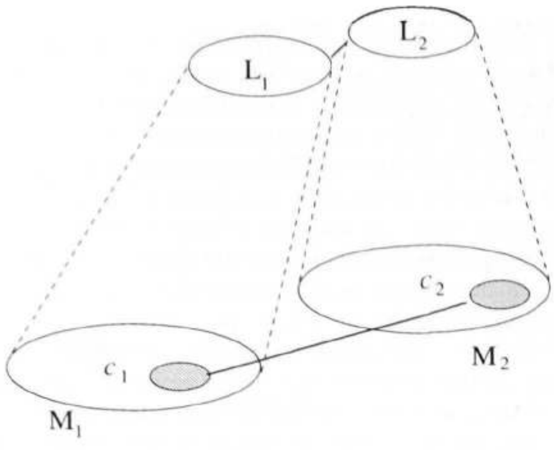

Figure 1: Local Models Semantics

(Giunchiglia & Ghidini, 1998) provides an intuitive semantics for MC Systems, called *local models semantics* (LMS). A detailed presentation of LMS is beyond the scope of this paper. Here we present the core intuitions, mainly to show how locality and compatibility are captured by LMS. For the sake of simplicity, we discuss the case with only two contexts. The generalization to any enumerable set of contexts is straightforward.

Consider Figure 1. L_1, L_2 are two first order languages (think of them as the languages associated with two contexts c_1 and c_2). The circles labelled M_1, M_2 at the bottom of the figure are the class of all first order models for the languages L_1, L_2 taken in isolation. Each $m \in M_i$ (for $i = 1, 2$) is called a *local model* for L_i. The basic idea of LMS is that there may exist some constraints on the local models of a language when this language is "embedded" in a MC System. These constraints (called *compatibility relations*) are the semantical counterpart of bridge rules. Syntactically, a bridge rule allows us to derive in a contexts more formulae than we would derive using only local inference rules. Semantically, this means that the set of (admissible) local models is smaller than it would be without any bridge rule. Therefore the semantical effect of a bridge rule is to cut away local models of a language. A *context* – from a semantical point of view – is thought of as the set of local models which not only satisfy a set of local axioms (the Ω of the definition), but also a set of further constraints which depend on the relations with the local models of other languages. In Figure 1, the two grey areas within the circles labelled M_1, M_2 (respectively c_1, c_2) represents two contexts, namely the result of restricting the sets M_1, M_2 (*i.e.* the models of the two languages taken in isolation) to the set of models that "survived" the further constraints.

Formally, let $\overline{m}_1 \subseteq M_1$ and $\overline{m}_2 \subseteq M_2$ be two sets of local models for L_1 and L_2. The pair $\langle \overline{m}_1, \overline{m}_2 \rangle$ is called a compatibility pair. A *compatibility relation* $\mathbf{C}$ is a set of compatibility pairs, namely:

$$\mathbf{C} \subseteq 2^{M_1} \times 2^{M_1}$$

A *model* for a MC System is a non-empty compatibility relation such that the pair $\langle \emptyset, \emptyset \rangle$ does not belong to the relation (this is for technical reasons that we do not need to discuss here). We say that a set of local models $\overline{m}_i$ *locally satifies* a formula $c_i : \Phi$ (in symbols, $\overline{m}_i \models_{loc} i : \Phi$) if and only if $\forall m \in \overline{m}_i (m \models_{CL} \phi)$ ($\models_{CL}$ is the symbol for classical satisfaction). A model $\mathbf{C}$ *satisfies* a formula $c_i : \Phi$ if and only if for any compatibility pair $\langle \overline{m}_1, \overline{m}_2 \rangle \in \mathbf{C}$, $\overline{m}_1 \models_{loc} c_1 : \Phi$ (for $i = 1$) and $\overline{m}_2 \models_{loc} c_2 : \Phi$ (for $i = 2$).

Proof-theoretically, locality is captured by the application of the rules Δ_i associated with a context c_i, and compatibility is captured by the application of bridge rules. Model-theoretically, locality is captured by the fact that the notion of satisfiability is local (the satisfiability of a (labelled) formula is given in terms of the local satisfiability of the formula with respect to its context, and the structures we consider to test local satisfiability are local: contexts have their own, generally different, domains of interpretation, sets of relations, and sets of functions); compatibility is captured by the idea of forcing relations over the set of local models defining each context (in some sense, we force local models to agree up to a certain extent).

Let us reconsider the first example we gave at the end of the section on locality and compatibility. c_1 was meant to describe John's beliefs on January 31, 1998 and c_2 to describe John's beliefs on February 1st, 1998. The intuitive relation between the two contexts c_1 and c_2 can be formalized in a MC system as follows:

- proof-theoretically, we can provide a bridge rule of the form:

$$\frac{c_1 : \Phi[\textit{Today}]}{c_2 : \Phi[\textit{Yesterday}]}$$

  (where $\Phi[t]$ stands for any formula whose truth is relative to a time t). It allows us to infer *Sunny(Yesterday)* in c_2 whenever we can infer *Sunny(Today)* in c_1;

- model-theoretically, we impose the following constraint over local models of c_1 and c_2: two local models m_1 and m_2 (of c_1 and c_2 respectively) are compatible iff whenever the first locally satisfies a sentence of the form $\Phi[Today]$, then the second locally satisfies a sentence of the form $\Phi[Yesterday]$.

Contexts in Mental Representation

Our intuition is that many important mechanisms which were proposed in mental represention can be quite directly mapped onto the principles of locality and compatibility, and therefore formalized as a suitable class of MC systems. Here we consider three of these mechanisms: the existence of a (logical) structure in an agent's mental contents, the (dynamic) generation of reasoning spaces where problems are solved, and the capability to carry on reasoning processes that cut across different reasoning spaces.

Structuring mental contents

Many authors believe that mental contents of an agent are better thought of as a collection of relatively small sets of facts rather than as a unique, huge and unstructured repository. The reason for such a structuring of mental contents is that it provides an economical and intuitive explanation of many reasoning processes which are not easily accounted for in a different conceptual framework. These sets of facts have been given different names: mental spaces (Fauconnier, 1985), spaces (Dinsmore, 1991), micro-theories (Lenat & Guha, 1990; Guha, 1991), contexts (Guha, 1991; McCarthy, 1993). Dinsmore writes that "each space represents some logically coherent situation or potential reality, in which various propositions are treated as true, objects are assumed to exist, and relations between objects are supposed to hold". Guha defines a micro-theory as "a theory of some topic, e.g. a theory of mechanics, a theory of the weather in the winter, a theory of what to look for when buying cars, etc." We believe that the notion of context we proposed in the previous sections can be used to formalize all these intuitions, since each space (mental space, micro-theory, ...) can be thought of as a partial and approximate theory (*i.e.* as a context).

Reasoning spaces

There is much evidence that many reasoning processes are highly contextual. The underlying intuition is that real agents do not use all they know in order to solve a problem, but rather a small subset of it. For instance, when playing chess, we are unlikely to consider what we know about the Italian cuisine or about Sherlock Holmes, unless we have good reasons to believe that these topics are relevant in order to win the game. In problem solving, we use only a subset of the facts that are potentially available (both from our memory and from the external environment), *i.e.*, the set of facts that we assume to be relevant to solve the problem at hand[2]. Sperber and Wilson, in their well-known book about relevance (Sperber & Wilson, 1986), discuss a lot of nice examples to show that the interpretation of a speaker's utterance happens in a context. In their terminology, a context is a collection of sentences expressing "expectations about the future, scientific hypotheses or religious beliefs, anecdotal memories, general cultural assumptions, beliefs about the mental state of the speaker". It is also worth mentioning that there are philosophers who adopt a similar perspective; here we only cite Davidson's explanation of the communication process in (Davidson, 1986). (Dinsmore, 1991) calls "parochial reasoning" any form of reasoning which involves only the resource available in a given space. (Guha, 1991; McCarthy, 1993; Buvac & Mason, 1993) use the metaphor of "entering" a context to say that reasoning can be circumscribed only to the set of facts which are true in a context. Behind all these concepts, the unifying idea seems to be that reasoning is often localized to a portion of an agent's global state. As such, a reasoning space can be thought of as a context where only local facts and rules are used.

Reasoning across different spaces

There are many senses in which reasoning processes may involve information belonging to different (mental) spaces. A first possibility is that a fact (a set of facts) is imported from a context to another one in order to solve a problem. A nice example is the Glasgow-London-Moscow problem (McCarthy, 1991). Suppose an agent is planning a trip from Glasgow to Moscow via London. The plan will include flight schedules, airtickets, airports, and so on. Now suppose we ask the agent: "What if your pants are stolen at the airport toilet?" Generally a human being has no problem to recognize that this is an (unexpected) obstacle to his plan. McCarthy's point is that it is unreasonable to assume that facts about clothes and social norms are by default included in a contenxt concerning air travelling. Therefore his conclusion is that they must be lifted from another context on the basis of some relation of generality over contexts (*e.g.* the context of social rules is more general than the context of air travelling because the second is a special case of a social action)[3] The idea that the contents of one space can be inherited by another space is discussed also in Dinsmore's work on partitioned representations.

Inheritance is just one possible relation between contexts (spaces, micro-theories). Other important relations are *specialization* (the contents of a space may specialize the contents of another context to a particular place, time, agent, etc.) and *reflection* (a space can be a sort of meta-level description of the contents of another space). The general mechanism which links the contents of a space to the contents of another space is analyzed in different terms. (Guha, 1991; McCarthy, 1993) call it *lifting* and define it as follows: "Given a pair of contexts C_1 and C_2 and a formula F_1 in C_1 we would like to determine what an equivalent formula in C_2 would be; *i.e.*, we are interested in computing a formula F_2 which in C_2 would 'state the same thing' in C_2 as F_1 in C_1" (this definition emphasizes the fact that lifting a fact from a context to another in general requires a "translation" in order to preserve its intended meaning). In the logic of PR, Dinsmore introduces two notions of context to explain the relation between different spaces. On one side, *primary contexts* are defined as "representations of functions that map the satisfaction of a proposition in one space onto the satisfaction of a (more complex) proposition into another space"[4] On the other side, *secondary contexts* are defined as "a kind of mapping from the content of one space to the contents of another that is consequence of the semantics of the primary contexts involved". The idea is that the mapping strictly depends upon the intuitive meaning of the contents of the spaces.

In our approach, all these examples can be given a simple interpretation in terms of bridge rules (semantically, in terms of compatibility relations). Any possible relation between contexts (inheritance, specialization, reflection, ...) corresponds to a different (set of) bridge rule(s). For instance, a simple form of inheritance can be expressed via the bridge rule:

$$\frac{c_1 : P}{c_2 : P}$$

The intuitive meaning is that, for any L_1-formula P that can be proved in c_1, there is a L_2-formula P that can be proved in c_2[5] Reflection can be expressed via the following

[2]It goes without saying that this process of "focussing" has evident efficiency advantages, but also that it can be used to explain many reasoning failures (since some relevant facts can be disregarded).

[3]A different approach, based on MC Systems, can be found in (Bouquet & Giunchiglia, 1995).

[4]In short, the idea is the following. Suppose a space **S1** represents John's beliefs and contains the sentence "Mary is nice" Then the statement "John believes that [[**S1**]]" is the primary context of **S1**. From "Mary is nice" in **S1**, the primary context allows us to infer the (more complex) sentence "John believes that Mary is nice"

[5]Notice that, because of locality, we cannot say that the formula P is true in both contexts, since P in L_1 and P in L_2 are not the same formula. We will come back to this problem in the next section.

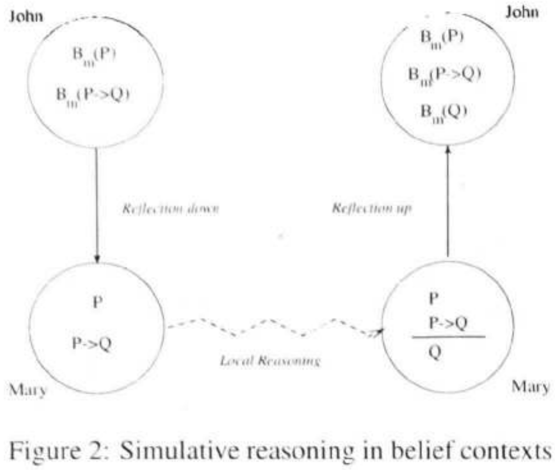

Figure 2: Simulative reasoning in belief contexts

two rules (called reflection up and down respectively):

$$\frac{c_1 : P}{c_2 : \bullet(P)} \qquad \frac{c_2 : \bullet(P)}{c_1 : P}$$

where the symbol '•' represents any possible predicate on which we operate the reflection (common examples are: Th – is a theorem; $Prov$ – is provable; B_i – is a belief of the agent i). Figure 2 shows an application of reflection rules in order to perform simulative reasoning on beliefs. The circle labelled with "John" on the left represents some of the beliefs that John ascribes to Mary. In order to simulate Mary's reasoning on her beliefs, John first reflects them down in a belief context (the circle labelled "Mary" in the picture) representing Mary's beliefs from John's perspective; second, performs local reasoning within this context (basically, a trivial application of Modus Ponens); finally reflects up the result in his own belief context[6]. Notice that the formula $B_m(Q)$ is derived in the context "John" because of the compatibility relation formalized by the two reflection rules. In other words, all local models of the context "John" which satisfy the formulae $B_m(P)$ and $B_m(P \rightarrow Q)$ and do not satisfy the formula $B_m(Q)$ are not compatible with the local models of the context "Mary" which satisfy the formulae P and $P \rightarrow Q$.

Comparisons with other frameworks

The goal of this conclusive section is to emphasize the differences with other proposed frameworks both from an intuitive and a technical point of view.

An important feature of MC Systems is that they have a flat architecture. Metaphorically, they can be viewed as a flat collection of theories possibly linked via a set of logical constraints. This is very different from what happens with Dinsmore's logic of PR and with McCarthy's logic of context. In the first case, we have a top down architecture, at the bottom of which there is a special space called **base**. The idea is that the asserttions of any space can be given a content if and only if they can be "translated" into an assertion of the space **base** via the rule of context climbing; otherwise, PR doesn't provide an interpretation for the contents of that space. McCarthy's logic of context has a bottom-up architecture, namely it has the structure of a onion (actually, several onions) with infinitely many layers. The idea is that any context can be always transcended via the operation of leaving a context: if the formula Φ is true in the context c, then there is always a context c_1 such that the formula $ist(c_1, \Phi)$ is true in it; and this can go on indefinitely. The intuition is that, since we can never make a complete list of the implicit assumptions of a context, there is always a more general context in which some new assumption is taken into account. The problem with these two approaches is that they do not respect the principle of locality: Dinsmore reduces the truth of an assertion in a space to the truth of a (more complex) assertion in another space (*i.e.* the space **base**); McCarthy's logic of context (at least in the form it is formalized in (Buvac & Mason, 1993)) introduces modal formulae whose truth value (in a context) depends on the truth value of another formula in a different context (and therefore there is not a clear notion of local satisfiability). It is worth noting, however, that MC systems allow us to "simulate" the hierarchical structure of McCarthy's logic. Indeed, we can define a MC system whose bridge rules are reflection rules of the form:

$$\frac{c_i : P}{c_{i+1} : ist(c_i, P)}$$

whose intuitive meaning is that for any index i, the formula $ist(c_i, P)$ is derivable in c_{i+1} iff the formula P is derivable in c_i. The main difference is that this hierarchy of embedded contexts is "hardwired" in McCarthy's logic, whereas in our framework it is just one possible architecture which represents a special case of compatibility involving an infinite chain of contexts such that the context c_{i+1} is a sort of metatheory of the context c_i.

As far as we know, MC systems are the only formal framework that explicitly allows for multiple distinct languages. Dinsmore, Guha, McCarthy, Buvac all introduce many examples which seem to require such a feature. However, when they turn from the examples to the formal systems, things are different. Dinsmore's logic of PR and Buvac's logic of context explicitly start with a single language, so that the set of well-formed formulae is the same in every space (context). Guha's logic is the only exception, but his distinction between grammaticality and meaningfulness results in a definition of language (starting from a single vocabulary) that is quite cumbersome. Beside the localization of expressiveness, the conceptual advantage of having multiple languages is that we do not have the problem of defining the different meanings of the same symbol in different contexts (single-language approach), but we can just specify the relations between the interpretation of symbols belonging to the language of different contexts. In other words, the fact that two formulae belonging to the languages of different contexts happen to have the same syntactic form does not mean in principle that their meaning is related in any way. If such a relation exists, then we must state it explicitly by adding a suitable bridge rule (or axiom) in the system.

MC Systems are the only context-based framework which allows for different logics in different contexts. In all other proposed frameworks, the logic of the system is global, namely we cannot use different sets of rules in different con-

[6] The two circles labelled "John" are not two different belief contexts, but the same context after some reasoning has been performed. Analogously for the context "Mary"

texts. We have already discussed the example of air travelling vs. looking for a phone number. Another example concerns simulative reasoning. Consider again Figure 2. If John had good reasons to assume that Mary cannot use Modus Ponens, then such a rule would not be included among the inference rules of the context "Mary" (the context where John simulates Mary's reasoning process). In our opinion, this should be an essential feature of any framework for mental representation.

Compatibility allows us to model most of the examples of lifting proposed by Guha, McCarthy and his group, and the examples of partitioned reasoning proposed by Dinsmore. We stress the fact that compatibility is conceptually very different from "lifting" as defined by Guha and McCarthy. In MC systems, we do not have the notion of "saying the same thing" in two different contexts. It is interesting to notice that this notion seems in contraddiction with the intuitions that Guha and McCarthy themselves describe in their work, namely that we can never have complete knowledge of a context. If so, it is not clear how we could then know that two sentences in different context express the "same" content, if not for trivial cases. Compatibility is a way out of this difficulty, since we only require that the intended meaning of two formule Φ and Ψ belonging to contexts c_1 and c_2 respectively is such that if Φ is locally satisfied by a set of local models $\overline{m}_1$, then Ψ must be locally satisfied by all the sets of local models $\overline{m}_2$ which stands in a compatibility relation with $\overline{m}_1$. We believe this is the best that we can do with contextual information as far as we accept that contexts are only partially known.

Acknowledgments

The authors would like to thank the members of *Mechanized Reasoning Group* at ITC-IRST (Trento), DISA (University of Trento), and DIST (University of Genoa). In particular, Massimo Benerecetti, Chiara Ghidini and Luciano Serafini gave the authors very useful feedback. This work is part of the MRG project on *Distributed Representations and Systems*.

References

Bouquet, P., Castelfranchi, C., Giunchiglia, F., & Kokinov, B. (Eds.) (1997). *ECCS'97 Workshop on Context*, Manchester (UK). AISB. Working Notes.

Bouquet, P. & Giunchiglia, F. (1995). Reasoning about Theory Adequacy. A New Solution to the Qualification Problem. *Fundamenta Informaticae*, 23(2–4):247–262.

Brezillion, P. (Ed.) (1997). *First International and Interdisciplinary Conference on Modeling and Using Context*, Rio de Janeiro (Brazil).

Brezillon, P. (Ed.) (1993). *Using Knowledge in its Context*, Chambery (France). IJCAI.

Brezillon, P. & Abu-Hakima, S. (Eds.) (1995). *Modelling Context in Knowledge Representation and Reasoning*, Montreal, Canada. IJCAI.

Buvac, S. & Ivanska, L. (Eds.) (1997). *Context in Knowledge Representation and Natural Language*, AAAI-97 Fall Symposium Series, MIT, Cambridge (MA). AAAI. Working Notes.

Buvac, S. & Mason, I. A. (1993). Propositional logic of context. *Proc. of the 11th National Conference on Artificial Intelligence*.

Davidson, D. (1986). A Nice Derangement of Epitaphs. In Lepore, E. (Ed.), *Truth and Interpretation. Perspectives on the Philosophy of Donald Davidson*. Basil Blackwell.

Dinsmore, J. (1991). *Partitioned Representations*. Kluwer Academic Publisher.

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

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.

Giunchiglia, F. & Ghidini, C. (1998). Local models semantics, or contextual reasoning = locality + compatibility. In *Sixth International Conference on Principles of Knowledge Representation and Reasoning (KR'98)*, Trento (Italy).

Giunchiglia, F. & Serafini, L. (1994). Multilanguage hierarchical logics (or: how we can do without modal logics). *Artificial Intelligence*, 65:29–70.

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

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

Kokinov, B. (1995). A Dynamic Approach to Context Modelling. In Brezillon, P. & Abu-Hakima, S. (Eds.), *Working Notes of the IJCAI-95 Workshop on "Modelling Context in Knowledge Representation and Reasoning"*, Montreal (Canada).

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

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

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. (1991). Overcoming an Unexpected Obstacle. Unpublished.

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

Prawitz, D. (1965). *Natural Deduction - A proof theoretical study*. Almquist and Wiksell, Stockholm.

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