2

Formalizing Context (Expanded Notes)

John McCarthy and Saša Buvač

These notes discuss formalizing contexts as *first class objects*. The basic relations are

$ist(c, p)$ meaning that the proposition p is true in the context c, and

$value(c, e)$ designating the value of the term e in the context c.

Besides these there are *lifting formulas* that relate the propositions and terms in subcontexts to possibly more general propositions and terms in the outer context. Subcontexts are often specialized with regard to time, place and terminology.

Introducing contexts as formal objects will permit axiomatizations in limited contexts to be expanded to *transcend* the original limitations. This seems necessary to provide AI programs using logic with certain capabilities that human fact representation and human reasoning possess. Fully implementing transcendence seems to require further extensions to mathematical logic, i.e. beyond the nonmonotonic inference methods first invented in AI and now studied as a new domain of logic.

1 Introduction

These notes contain some of the reasoning behind the proposals of Mc-Carthy 1987 to introduce contexts as formal objects. The present proposals are incomplete and tentative. In particular the formulas are not what we will eventually want, and we will feel free to use formulas in discussions of different applications that aren't always compatible with each other. This is an expanded and revised version of McCarthy 1993. An earlier version of this paper is the Stanford University Technical Note STAN-CS-TN-94-13. The current version contains new sections §7 and §13, as well as updated bibliographical remarks. Some of the results in this paper have been previously also published in one of the following: Buvač and McCarthy 1996, Buvač 1996a, Buvač 1996b.

Our object is to introduce contexts as abstract mathematical entities with properties useful in artificial intelligence. Our attitude is therefore a computer science or engineering attitude. If one takes a psychological or philosophical attitude, one can examine the phenomenon of contextual dependence of an utterance or a belief. However, it seems to us unlikely that this study will result in a unique conclusion about *what context is.* Instead, as is usual in AI, various notions will be found useful.

One major AI goal of this formalization is to allow simple axioms for common sense phenomena, e.g. axioms for static blocks world situations, to be *lifted* to contexts involving fewer assumptions, e.g. to contexts in which situations change. This is necessary if the axioms are to be included in general common sense databases that can be used by any programs needing to know about the phenomenon covered but which may be concerned with other matters as well. Rules for lifting are described in section 4 and an example is given.

A second goal is to treat the context associated with a particular circumstance, e.g. the context of a conversation in which terms have particular meanings that they wouldn't have in the language in general.

The most ambitious goal is to make AI systems which are never permanently stuck with the concepts they use at a given time because they can always *transcend* the context they are in—if they are smart enough or are told how to do so. To this end, formulas $ist(c, p)$ are always considered as themselves asserted within a context, i.e. we have something like $ist(c', ist(c, p))$. The regress is infinite, but we will show that it is harmless.

The main formulas are sentences of the form

$$c' : \quad ist(c, p),$$

which are to be taken as assertions that the proposition p is true in the context c, itself asserted in an outer context c'. (We have adopted

Guha's Guha 1991 notation rather than that of McCarthy 1987, because he built his into Cyc, and it was easy for us to change ours.) For now, propositions may be identified with sentences in English or in various logical languages, but we may later take them in the sense of McCarthy 1979b as abstractions with possibly different identity conditions. We will use both logical sentences and English sentences in the examples, according to whichever is more convenient.

Contexts are abstract objects. We don't offer a definition, but we will offer some examples. Some contexts will be *rich* objects, like situations in situation calculus. For example, the context associated with a conversation is rich; we cannot list all the common assumptions of the participants. Thus we don't purport to describe such contexts completely; we only say something about them. On the other hand, the contexts associated with certain microtheories are poor and can be completely described.

Here are some examples.

$c0$: $ist(context\text{-}of(\text{“Sherlock Holmes stories”}),$

     $\text{“Holmes is a detective”})$

asserts that it is true in the context of the Sherlock Holmes stories that Holmes is a detective. We use English quotations here, because the formal notation is still undecided. Here $c0$ is considered to be an *outer context*. In the context $context\text{-}of(\text{“Sherlock Holmes stories”})$, Holmes's mother's maiden name does not have a value. We also have

$c0$: $ist(context\text{-}of(\text{“U.S. legal history”}),$

     $\text{“Holmes is a Supreme Court Justice”}).$

Since the outer context is taken to be the same as above, we will omit it in subsequent formulas until it becomes relevant again. In this context, Holmes's mother's maiden name has a value, namely Jackson, and it would still have that value even if no-one today knew it.

$ist(c1, at(jmc, Stanford))$ is the assertion that John McCarthy is at Stanford University in a context in which it is given that jmc stands for the first author of this paper and that $Stanford$ stands for Stanford University. The context $c1$ may be one in which the symbol at is taken in the sense of being regularly at a place, rather than meaning momentarily at the place. In another context $c2$, $at(jmc, Stanford)$ may mean physical presence at Stanford at a certain instant. Programs based on the theory should use the appropriate meaning automatically.

Besides the sentence $ist(c, p)$, we also want the term $value(c, term)$ where $term$ is a term. For example, we may need $value(c, time)$, when c is a context that has a time, e.g. a context usable for making asser-

tions about a particular situation. The interpretation of $value(c, term)$ involves a problem that doesn't arise with $ist(c, p)$. Namely, the space in which terms take values may itself be context dependent. However, many applications will not require this generality and will allow the domain of terms to be regarded as fixed.

Here's another example of the value of a term depending on context:

$c0$: $value(context\text{-}of(\text{"Sherlock Holmes stories"}),$

"number of Holmes's wives") $= 0$

whereas

$c0$: $value(context\text{-}of(\text{"U.S. legal history"}),$

"number of Holmes's wives") $= 1.$

We can consider $setof\text{-}wives(Holmes)$ as a term for which the set of possible values depends on context. In the case of the Supreme Court justice, the set consists of real women, whereas in the Sherlock Holmes case, it consists of fictitious women.

The remainder of this paper is organized as follows. In §2 we give examples of some elementary relations among contexts. The basic operations of contextual reasoning, entering and exiting contexts, are introduced in §3. In §4 we focus on *lifting axioms*—axioms relating what is true in one context based on what is true in another context. Building on the basic notions of entering/exiting contexts and lifting axioms, §5 shows how contexts can be used to reason in the style of natural deduction. To illustrate short term applicability of contexts, §6 demonstrates how the context formalism aids in the integration of databases which were not originally intended to be used together. These techniques are not specific to databases and can be applied in integrating other objects, such as plans; see §7. In §8 we treat contexts associated with particular circumstances, namely those that come up in a conversation. The transcending of the outer context of a system, as is discussed in §9, might result in AI programs which are never permanently stuck with the concepts they use at a particular time. In §10, we argue that all sentences will always be context dependent, and thus it is not possible to define an absolute outermost context. Returning to applications, in §11 we sketch how contexts can be used to represent mental states and revise the beliefs of an agent. We conclude with a some remarks and conclusions in §12 and §13.

2 Relations among Contexts

There are many useful relations among contexts and also context valued functions. Here are some.

1. *specialize-time*(t, c) is a context related to c in which the time is specialized to have the value t. We may have the relation

$$c0: \quad ist(\textit{specialize-time}(t, c), \textit{at}(\textit{jmc}, \textit{Stanford})) \equiv$$

$$ist(c, \textit{at-time}(t, \textit{at}(\textit{jmc}, \textit{Stanford}))).$$

Here *at-time*(t, p) is the assertion that the proposition p holds at time t. We call this a *lifting* relation. It may be convenient to write *at-time*$(t, \textit{foo}(x, y, z))$ rather than $\textit{foo}(x, y, z, t)$, because this lets us drop t in certain contexts. Many expressions are also better represented using modifiers expressed by functions rather than by using predicates and functions with many arguments. Actions give immediate examples, e.g. *slowly*$(\textit{on-foot}(go))$ rather than $go(\textit{on-foot}, \textit{slowly})$.

Instead of using the function *specialize-time*, it may be convenient to use a predicate *specializes-time* and an axiom

$$c0: \quad \textit{specializes-time}(t, c1, c2) \wedge ist(c1, p) \supset ist(c2, \textit{at-time}(t, p)).$$

This would permit different contexts $c1$ all of which specialize $c2$ to a particular time.

There are also relations concerned with specializing places and with specializing speakers and hearers. Such relations permit lifting sentences containing pronouns to contexts not presuming specific places and persons.

2. If q is a proposition and c is a context, then *assuming*(p, c) is another context like c in which p is assumed, where "assumed" is taken in the natural deduction sense. We investigate this further in §5.

3. There is a general relation *specializes* between contexts. We say *specializes*$(c1, c2)$ when $c2$ involves no more assumptions than $c1$. We have nonmonotonic relations

$$\textit{specializes}(c1, c2) \wedge \neg ab1(p, c1, c2) \wedge ist(c1, p) \supset ist(c2, p),$$

and

$$\textit{specializes}(c1, c2) \wedge \neg ab2(p, c1, c2) \wedge ist(c2, p) \supset ist(c1, p).$$

This gives nonmonotonic inheritance of *ist* both from the subcontext to the supercontext and vice versa. More useful is the case when the sentences must change when lifted. Then we need to state that and

every *proposition* meaningful in $c1$ is translatable into one meaningful in $c2$. See §4 for an example.

4. A major set of relations that need to be expressed are those between the context of a particular conversation and a subsequent written report about the situation in which the conversation took place. References to persons and objects are *decontextualized* in the report, and sentences like those given above can be used to express their relations.

5. Consider a wire with a signal on it which may have the value 0 or 1. We can associate a context with this wire that depends on time. Call it $c_{wire117}(t)$. Suppose at time 331, the value of this signal is 0. We can write this

$$ist(c_{wire117}(331), signal = 0).$$

Suppose the meaning of the signal is that the door of the microwave oven is open or closed according to whether the signal on *wire*117 is 0 or 1. We can then write the lifting relation

$$(\forall\ t)(ist(c_{wire117}(t), signal = 0) \equiv door\text{-}open(t).$$

The idea is that we can introduce contexts associated with particular parts of a circuit or other system, each with its special language, and lift sentences from this context to sentences meaningful for the system as a whole.

3 Entering and Exiting Contexts

Suppose we have the formula $c0 :\ ist(c, p)$. We can then *enter* the context c and infer the formula $c : p$. Conversely, if we have the formula $c :\ p$ we can infer $c0 :\ ist(c, p)$ by *exiting* the context c. We don't always want to be explicit about the sequence of all the contexts that were entered, but the logic needs to be such that the system always exits into the context it was in before entering. The enter and exit operations can be thought of as the push and pop operations on a stack. In the logic presented in Buvač et al. 1995 the sequence of contexts that has been entered is always explicitly stated.

We can regard $ist(c, p)$ as analogous to $c \supset p$, and the operation of entering c as analogous to *assuming* c in a system of *natural deduction* as invented by Gentzen and described in many logic texts. Indeed a context is a generalization of a collection of assumptions, but there are important differences. For example, contexts contain linguistic *assumptions* as well as declarative and a context may correspond to an infinite and only partially known collection of assumptions. Moreover, because

relations among contexts are expressed as sentences in the language, $ist(c, p)$ allows inferences within the language that could only be done at the meta-level of the usual natural deduction systems.

There are various ways of handling the reasoning step of entering a context. The way most analogous to the usual natural deduction systems is to have an operation *enter c*. Having done this, one could then write any p for which one already had $ist(c, p)$. However, it seems more convenient in an interactive theorem proving to use the style of Jussi Ketonen's EKL interactive theorem prover Ketonen and Weening 1984. In the style of that system, if one had $ist(c, p)$, one could immediately write p, and the system would keep track of the dependence on c. To avoid ambiguity as to where an occurrence of $ist(\ , p)$ came from, one might have to refer to a line number in the derivation. Having obtained p by entering c and then inferring some sentence q, one can *leave c* and get $ist(c, q)$. In natural deduction, this would be called discharging the assumption c.

Human natural language risks ambiguity by not always specifying such assumptions, relying on the hearer or reader to guess what contexts makes sense. The hearer employs a *principle of charity* and chooses an interpretation that assumes the speaker is making sense. In AI usage we probably don't usually want computers to make assertions that depend on principles of charity for their interpretation.

We are presently doubtful that the reasoning we will want our programs to do on their own will correspond closely to using an interactive theorem prover. Therefore, it isn't clear whether the above ideas for implementing entering and leaving contexts will be what we want.

Sentences of the form $ist(c, p)$ can themselves be true in contexts, e.g. we can have $ist(c0, ist(c1, p))$. In this draft, we will ignore the fact that if we want to stay in first order logic, we should reify assertions and write something like $ist(c0, Ist(c1, p))$, where $Ist(c, p)$ is a term rather than a wff. Actually the same problem arises for p itself; the occurrence of p in $ist(c, p)$ might have to be syntactically distinct from the occurrence of p standing by itself. Alternatively to reifying assertions we could use a modal logic; this approach is investigated in Shoham 1991, Buvač 1996a, Amati and Pirri 1997, van Benthem 1997.

4 Lifting Axioms

Lifting axioms are axioms which relate the truth in one context to the truth in another context. Lifting is the process of inferring what is true in one context based on what is true in another context by the means of lifting axioms. We treat lifting as an informal notion in the sense that

we never introduce a lifting operator. In this section we give an example of lifting. See Buvač and Fikes 1995 for more examples.

Consider a context *above-theory*, which expresses a static theory of the blocks world predicates *on* and *above*. In reasoning about the predicates themselves it is convenient not to make them depend on situations or on a time parameter. However, we need to *lift* the results of *above-theory* to outer contexts that do involve situations or times.

To describe *above-theory*, we may write informally

(1)   *above-theory* : $(\forall xy)(on(x,y) \supset above(x,y))$

(2)   *above-theory* : $(\forall xyz)(above(x,y) \wedge above(y,z) \supset above(x,z))$

<div align="center">etc.</div>

which stands for

(3) $c0$: $ist(above\text{-}theory, (\forall xy)(on(x,y) \supset above(x,y)))$

<div align="center">etc.</div>

Constant $c0$ denotes an outer context. Section §9 has more about $c0$. In the following formulas, we put the context in which the formula is true to the left followed by a colon.

We want to use the *above-theory* in a context *blocks* which contains the theory of blocks world expressed using situation calculus. (We assume that situations are a disjoint sort, and that the variable s ranges over the situation sort.) In the context *blocks* predicates *on* and *above* have a third argument denoting a situation. Thus the context *blocks* needs to relate its three-argument predicates $on(x,y,s)$ and $above(x,y,s)$ to two-argument predicates $on(x,y)$ and $above(x,y)$ of the *above-theory* context. This is done by introducing a context of a particular situation, *spec-sit*(s). A context *spec-sit*(s) is associated with each situation s, such that

(4)   *blocks* : $(\forall xys)(on(x,y,s) \equiv ist(spec\text{-}sit(s), on(x,y)))$,

(5)   *blocks* : $(\forall xys)(above(x,y,s) \equiv ist(spec\text{-}sit(s), above(x,y)))$,

<div align="center">etc.</div>

In order to get relations between $on(x,y,s)$ and $above(x,y,s)$, we will now import *above-theory* into the *blocks* context. The importation of *above-theory* is expressed by the axiom

(6) $c0$: $(\forall p)ist(above\text{-}theory, p) \supset ist(blocks, (\forall s)(ist(spec\text{-}sit(s), p)))$,

asserting that the facts of *above-theory* all hold in the contexts associated with every situation. The following relation between $on(x,y,s)$ and $above(x,y,s)$ follows from the above axioms.

Theorem (above):

$$blocks: \quad (\forall sxy)(on(x,y,s) \supset above(x,y,s)).$$

The example given is so small that it would be simpler to give the relations among the three-argument predicates directly, but imagine that *above-theory* was much larger than is given here.

We proceed to derive the above theorem.

Proof (above): We begin by assuming

(7) $blocks: \quad on(x,y,s),$

asserting that block x is on block y in a specific situation s. Together with the universally instantiated form of the $\Rightarrow$ direction of formula 4 we get

(8) $blocks: \quad ist(spec\text{-}sit(s), on(x,y)).$

Now we enter *spec-sit(s)* and get

(9) $spec\text{-}sit(s): \quad on(x,y).$

From (3) and (6) we conclude

(10) $c0: \quad ist(blocks, (\forall s)ist(spec\text{-}sit(s), (\forall xy)on(x,y) \supset above(x,y))).$

Therefore, by entering *blocks* we have

(11) $blocks: \quad (\forall s)ist(spec\text{-}sit(s), (\forall xy)on(x,y) \supset above(x,y)).$

By universal instantiation it follows that

(12) $blocks: \quad ist(spec\text{-}sit(s), (\forall xy)on(x,y) \supset above(x,y)).$

Entering *spec-sit(s)* gives

(13) $spec\text{-}sit(s): \quad (\forall xy)on(x,y) \supset above(x,y).$

By logic, formulas 9 and 13 give

(14) $spec\text{-}sit(s): \quad above(x,y).$

We can now either continue reasoning in *spec-sit(s)* or exit *spec-sit(s)* and get

(15) $blocks: \quad ist(spec\text{-}sit(s), above(x,y)).$

Together with the universally instantiated form of the $\Leftarrow$ direction of formula 5 we get

(16) $blocks: \quad above(x,y,s).$

By the deduction theorem we can discharge the initial assumption to obtain

(17) $blocks: \quad on(x,y,s) \supset above(x,y,s).$

Finally, by universal generalization it follows that

(18) $\quad blocks : \quad (\forall sxy)on(x,y,s) \supset above(x,y,s).$

$\Box_{\mathbf{above}}$

In this derivation we used a function giving a context *spec-sit*(s) which depends on the situation parameter s. Contexts depending on parameters will surely present problems requiring more study.

Besides that, the careful reader of the derivation will wonder what system of logic permits the manipulations involved, especially the substitution of sentences for variables followed by the immediate use of the results of the substitution. There are various systems that can be used, e.g. quasi-quotation as used in the Lisp or KIF, use of back-quotes, or the ideas of McCarthy 1979b. Furthermore, the drafts of this paper have motivated a number of researchers to develop logics of context, in which (some version of) the above argument would be a derivation; these include Giunchiglia 1993, Nayak 1994, Buvač 1996a, Attardi and Simi 1998. However, at present we are more attached to the derivation than to any specific logical system.

As a further example, consider rules for lifting statements like those of section 1 to one in which we can express statements about Justice Holmes's opinion of the Sherlock Holmes stories.

5 Natural Deduction via Context

The formal theory of context can be used to represent inference and reason in the style of natural deduction. This requires lifting axioms (or lifting rules) to treat the context which a reasoning system is in as a formal object. If p is a sentence and we are in some context c, we define a new context *assuming*(c,p) so that it validates the following rules:

importation $\quad c : p \supset q \;\;\vdash\;\; assuming(c,p) : q$

discharge $\quad assuming(c,p) : q \;\;\vdash\;\; c : p \supset q$

Note that these rules can be replaced by lifting axioms. Thus **importation** is replaced by

(19) $\quad (\forall cpq)(ist(c,p \supset q) \supset ist(assuming(c,p),q))$

To make the presentation simpler we write them in the rule form. An interesting rule which can be derived from the above is

assumption $\quad \vdash assuming(c,p) : p$

In analogy to the restriction to the rule of $\forall$ introduction in formal systems of natural deduction, we will have to restrict the rule of universal generalization to ensure that the variable being generalized does not

occur free in any of the *assuming*(c, p) terms of the current context; see Prawitz 1965.

To illustrate the rules we now give a natural-deduction style proof of the **above** theorem, which was introduced in §4. This theorem involves the lifting of the theory of above into the context of situation calculus. The proof should be compared to the Hilbert style proof which was given in §4.

Proof (above):   We begin with the $\Rightarrow$ direction of formula 4

(20) *blocks* : $(\forall xys)(on(x, y, s) \supset ist(spec\text{-}sit(s), on(x, y)))$

It follows by universal instantiation that

(21) *blocks* : $on(x, y, s) \supset ist(spec\text{-}sit(s), on(x, y))$

By the **importation** rule we get

(22) *assuming*$(blocks, on(x, y, s))$: $ist(spec\text{-}sit(s), on(x, y))$

Therefore, by entering the *spec-sit*(s) context we get

(23) *spec-sit*(s) : $on(x, y)$

Now, from formulas 3 and 6 it follows that

(24) *c0* : $ist(blocks, (\forall s)ist(spec\text{-}sit(s), (\forall xy)(on(x, y) \supset above(x, y))))$

By entering *blocks* we get

(25) *blocks* : $(\forall s)ist(spec\text{-}sit(s), (\forall xy)(on(x, y) \supset above(x, y)))$

By instantiating the universal quantifier over situations we get

(26) *blocks* : $ist(spec\text{-}sit(s), (\forall xy)(on(x, y) \supset above(x, y)))$

Therefore, by propositional logic we have

(27) *blocks* : $on(x, y, s) \supset ist(spec\text{-}sit(s), (\forall xy)(on(x, y) \supset above(x, y)))$

Therefore, by the **importation** rule we get

(28) *assuming*$(blocks, on(x, y, s))$: $ist(spec\text{-}sit(s),$
$\qquad (\forall xy)(on(x, y) \supset above(x, y)))$

Now, by entering the *spec-sit*(s) context we get

(29) *spec-sit*(s) : $(\forall xy)(on(x, y) \supset above(x, y))$

By logic from formulas 23 and 29 it follows that

(30) *spec-sit*(s) : $above(x, y)$

By exiting the *spec-sit*(s) context we get

(31) *assuming*$(blocks, on(x, y, s))$: $ist(spec\text{-}sit(s), above(x, y))$

The $\Leftarrow$ direction of formula 5

(32) $blocks:$ $(\forall xys)ist(spec\text{-}sit(s), above(x, y)) \supset above(x, y, s)$

By propositional logic we have

(33) $blocks:$ $on(x, y, s) \supset$

$(\forall xys)ist(spec\text{-}sit(s), above(x, y)) \supset above(x, y, s)$

Together with the **importation** rule the above formula allows us to infer

(34) $assuming(blocks, on(x, y, s)):$

$(\forall sxy)ist(spec\text{-}sit(s), above(x, y)) \supset above(x, y, s)$

By logic from (31) and (34) we get

(35) $assuming(blocks, on(x, y, s)):$ $above(x, y, s)$

Using the rule **discharge** it follows that

(36) $blocks:$ $on(x, y, s) \supset above(x, y, s)$

Therefore, by universal generalization we obtain what was to be proved

(37) $blocks:$ $(\forall sxy)on(x, y, s) \supset above(x, y, s)$

$\square$**above**

In the above proof we have entered the context $assuming(c, p)$ in a number of instances. This creates an interesting example because it might not be obvious in which context the term $assuming(c, p)$ is to be interpreted. However, since the logic needs to keep track of which contexts were entered in the process of reasoning, the answer becomes obvious: the term $assuming(c, p)$ will be interpreted in the next outer context (see §3 for discussion on sequences of contexts).

We gave two treatments of the key argument: one in natural deduction (§4) and one Hilbert-style (above). This kind of proof transformation is logical routine, and any textbook on proof theory (say, the recent Troelstra and Schwichtenberg 1996), tells all about it.

5.1   Postponing Preconditions via *assuming*

We conclude by noting that the *assuming* function, as defined in this section, is also useful for formalizing a number of other phenomena. Examine a naive formalism for reasoning about action where the preconditions for flying are given by the formula

(38) $c:$ $have\text{-}ticket(x) \wedge clothed(x) \supset can\text{-}fly(x).$

In common sense reasoning we want the ability to *postpone* dealing with the precondition of being clothed. This can be done by considering a

context which assumes that one is clothed *assuming*(c, *clothed*(x)). By the **importation** rule and the formula 38 we get

(39) *assuming*(c, *clothed*(x)) :   *have-ticket*(x) ⊃ *can-fly*(x).

Thus in the context *assuming*(c, *clothed*(x)) we do not need to consider the precondition of being clothed in order to infer that one can fly.

Note that we are only developing an ontology for representing this phenomena, and are not dealing with pragmatic issues like which context a reasoning system will start in, and how the system will decide to consider a context making an additional assumption. In fact, from a pragmatic viewpoint the above process might need to be completely reversed. The reasoning system may realize that its current problem solving context c is making a particular assumption p that needs to be discharged. Then it will need to consider a context c' such that $c = assuming(c', p)$.

The *assuming* function is also needed for representing discourse. In §8 we show how it is used to handle replies to a query; in that section we call the *assuming* function *"reply"*.

See §7.2 for related examples in the planning domain.

6 Integrating Databases

We see the use of formalized contexts as one of the essential tools for reaching human level intelligence by logic based methods. However, formalized contexts have shorter term applications. In this section we deal with one short term application: we show how two data bases, which were not originally intended to be used together, can be integrated by lifting their contents into a wider context. We proceed with an example. For more practical issues involved in the task of integrating data and knowledge bases see Collet et al. 1991, Farquhar et al. 1995.

6.1 The GE, Navy, and Air Force Example

Here's a hypothetical example. Imagine that the Navy, the Air Force and General Electric have separately developed standards for databases containing facts about prices of jet engines and parts. But these standards are not the same. Suppose that associated with each item is a price. Suppose further

1. For GE, the price is a retail price not including spare parts.

2. For the Navy, the price is the Government's purchase price including spare parts.

3. For the Air Force, the price includes additional inventory costs. It includes spare parts but a different assortment than the Navy's.

Now suppose that associated with each database are many programs

that use the information. For example, General Electric can compute the cost of equipment packages taking into account discounts. The Navy can compute the economic ordering quantity for use when supplies get low.

Suppose now that some plan requires that unexpectedly a certain item made by General Electric is required in large quantity by both the Navy and the Air Force and deliveries and purchases from various General Electric warehouses have to be scheduled in co-ordination. The context in which the reasoning is done requires the *lifting* of various information from the contexts of the separate databases to the reasoning context. In the course of this lifting, the sentences representing the information are translated into new forms appropriate for the reasoning context.

6.2 A Simple Formalization

In this simple case, assume that the Air Force and Navy data bases need to be updated on the new General Electric prices. The GE database lists the list price, i.e. the price at which GE is selling the engine. The Navy database lists the price which Navy *will need to pay* for the engine and its assortment of spare parts, if it decides to use GE.

In order to reason with multiple databases, c_{ps}, an ad hoc context for reasoning about the particular problem, may be required. The problem solving context c_{ps} contains objects denoting the General Electric context c_{GE}, the Navy context c_N, and the Air Force context c_{AF}. This enables us to talk about facts which are contained in the corresponding databases. If for example the GE database contains a fact $price(FX\text{-}22\text{-}engine) = \$3600K$ then the sentence

$$ist(c_{GE}, price(FX\text{-}22\text{-}engine) = \$3600K)$$

is true in c_{ps}.

Different data bases might make different assumptions. For instance, prices of engines in some contexts might or might not include spare parts or warranties. We need the ability to represent this information in c_{ps}. Function *spares*, when given a product and a context, returns the spares which the given context assumes necessary and thus includes in the price of the product. For example, $spares(c_{NAVY}, x)$ is the set of spares that Navy assumes will be included in the price of the product x. Function *warranty*, when given a product and a context, returns the name of the warranty assumed for the product in the given context. For example, $warranty(c_{NAVY}, x)$ is the name of the warranty which Navy assumes is included in the price of the product x. In this note we are treating warranty in the same manner as we would treat spare parts or additional

optional features. It would be the responsibility of another formalization to "understand" the warranty and give axioms describing the exact obligations that GE has to its clients. Note that information about spares and warranties assumed by the Navy will probably not be contained in the Navy data base. (Otherwise, we would use $value(c_{\text{NAVY}}, spares(x))$ rather than $spares(c_{\text{NAVY}}, x)$ to refer to the spares that Navy assumes will be included in the price of the product x.) Rather, this information is kept in in some manual. But for these data bases to be used jointly, the spares information needs to be included; we assume that it is described declaratively in c_{ps}. Finally, the vocabulary of c_{ps} also has a function $GE\text{-}price$, which to every object assigns its corresponding price in dollars.

In the problem solving context c_{ps} we also represent the fact that GE lists engine prices without any spares, while Navy assumes spare parts to be included in the price of a product. This is done by lifting axioms, which define how the notion of price in different databases translates into the problem solving context:

(40) c_{ps} : $(\forall x)\,value(c_{\text{GE}}, price(x)) = GE\text{-}price(x)$

(41) c_{ps} : $(\forall x)\,value(c_{\text{NAVY}}, price(x)) = GE\text{-}price(x)+$

$\qquad GE\text{-}price(spares(c_{\text{NAVY}}, x))+$

$\qquad GE\text{-}price(warranty(c_{\text{NAVY}}, x))$

expressing that the price listed in the Navy data base is the price of the engine, some bag of spares, and the particular warranty that are assumed by the Navy.

(42) c_{ps} : $(\forall x)\,value(c_{\text{AF}}, price(x)) = f(x, GE\text{-}price(x),$

$\qquad GE\text{-}price(spares(c_{\text{AF}}, x)), GE\text{-}price(warranty(c_{\text{AF}}, x)))$

where f is some function which determines the total price of an item and spares, also taking into account the inventory cost. Note that f might not be precisely known, in which case we might decide to only give some approximate bounds on f.

Now we work out an example. Assume that we are given the prices as listed in the GE data base; i.e. the following formulas hold in c_{GE}:

(43) c_{GE} : $price(FX\text{-}22\text{-}engine) = \$3600K$

(44) c_{GE} : $price(FX\text{-}22\text{-}engine\text{-}fan\text{-}blades) = \$5K$

(45) c_{GE} : $price(FX\text{-}22\text{-}engine\text{-}two\text{-}year\text{-}warranty) = \$6K$

Information about spares and warranties will not be found in the c_{GE} data base and will probably require looking up in some manual or de-

scription of the the data base. We need to enter this information into the the problem solving context:

(46) c_{ps} : $spares(c_{\text{NAVY}}, FX\text{-}22\text{-}engine) = FX\text{-}22\text{-}engine\text{-}fan\text{-}blades$

(47) c_{ps} : $warranty(c_{\text{NAVY}}, FX\text{-}22\text{-}engine) =$
     $FX\text{-}22\text{-}engine\text{-}two\text{-}year\text{-}warranty$

Then we can compute the price of the FX-22 jet engine for the Navy. The following formula is a theorem, i.e. it follows from the above axioms.

Theorem (engine price):

c_{NAVY} : $price(FX\text{-}22\text{-}engine) = \$3611K$

In order to compute this price for the Air Force, the inventory cost given by function f would need to be known.

Proof (engine price):   First we exit the c_{GE} context thus rewriting formulas 43, 44, and 45 as

(48) c_{ps} : $value(c_{\text{GE}}, price(FX\text{-}22\text{-}engine)) = \$3600K$

(49) c_{ps} : $value(c_{\text{GE}}, price(FX\text{-}22\text{-}engine\text{-}fan\text{-}blades)) = \$5K$

(50) c_{ps} : $value(c_{\text{GE}}, price(FX\text{-}22\text{-}engine\text{-}two\text{-}year\text{-}warranty)) = \$6K$

From formulas 40 and 48 it follows that

(51) c_{ps} : $GE\text{-}price(FX\text{-}22\text{-}engine) = \$3600K$

From formulas 40 and 49 it follows that

(52) c_{ps} : $GE\text{-}price(FX\text{-}22\text{-}engine\text{-}fan\text{-}blades) = \$5K$

Therefore, using formula 46, we get

(53) c_{ps} : $GE\text{-}price(spares(c_{\text{NAVY}}, FX\text{-}22\text{-}engine)) = \$5K$

In a similar fashion, from formulas 40, 47 and 50 we can conclude that

(54) c_{ps} : $GE\text{-}price(warranty(c_{\text{NAVY}}, FX\text{-}22\text{-}engine)) = \$6K$

From formulas 51, 53, and 54 if follows that

(55) c_{ps} : $GE\text{-}price(FX\text{-}22\text{-}engine)+$
     $GE\text{-}price(spares(c_{\text{NAVY}}, FX\text{-}22\text{-}engine))+$
     $GE\text{-}price(warranty(c_{\text{NAVY}}, x)) = \$3611K$

Then, using formula 41 we can conclude that

(56) c_{ps} : $value(c_{\text{NAVY}}, price(FX\text{-}22\text{-}engine)) = \$3611K$

By entering c_{NAVY} we get

(57) $c_{\text{NAVY}} : \quad price(\textit{FX-22-engine}) = \$3611K$

$\square$**engine–price**

In the above proof we are assuming that all constants denote the same object in all contexts, i.e. that all constants are *rigid designators*. Consequently constants can be substituted for universally quantified variables by the universal instantiation rule. Generalizing the proof is straight forward if we drop this assumption.

6.3 Formalization for Bargaining

Now assume that the air force database contains the price air force *plans to pay* for a product, i.e. the price *included in the budget.* Like before, the GE database contain the list price, which will probably be higher than the air force budget price. This formalization is suited for use by some bargaining agents or programs. The bargaining agent for the air force will through negotiation attempt to convince the GE agent to lower the GE list price to the air force budget price (or some price that would be acceptable to the air force).

The bargaining agents will work in some problem solving context c_{ps}. This context contains constants denoting the various data bases which will be relevant to the bargaining; in our case these will be the General Electric context c_{GE}, and the Air Force context c_{AF}. Context c_{ps} contains functions which represent the budget price and the list price of a product. Function *manufacturer-price*, when given a context of a manufacturer and a product, returns the price at which the product is offered for sale by the manufacturer; functions *budget-price*, when given a context of a customer and a product, returns the price which the customer is willing to pay for the product. Like in the previous example, c_{ps} can represent the spares associated with an engine. Function *spares*, when given a product and an object, returns the spares which the given context assumes necessary and thus included in the price of the product.

The air force and GE will need to bargain in order to negotiate a price which is acceptable to both parties. However, since unlike GE, the air force assumes that the price will include a set of spare parts, the lifting axioms will be needed to adjust the prices in the two data bases to ensure that both the budget price and the list price pertain to the same package. The lifting axioms are:

(58) $c_{\text{ps}} : (\forall x) value(c_{\text{GE}}, price(x)) = \textit{manufacturer-price}(c_{\text{GE}}, x)$

(59) $c_{\text{ps}} : (\forall x) value(c_{\text{AF}}, price(x)) = \textit{budget-price}(c_{\text{AF}}, x) +$

$$budget\text{-}price(c_{AF}, spares(c_{AF}, x))$$

The lifting axioms will enable us to derive the *budget-price* and *manufacturer-price* prices in c_{ps}, both of which pertain to the engine only, excluding any spares. These can then be used by the bargaining programs to negotiate a price and administrate a sale.

Note again the difference between this formalization and the previous one. In the previous subsection the *price* function in both data bases referred to the price which was actually being paid for a product. Therefore, the lifting axioms were used to directly infer the price in one data base based on the price listed in another. In this example, on the other hand, given the list price the lifting axioms can not be used to work out the budget price. The lifting axioms simply ensure that both the list price and the budget price talk only about the engine, and do not implicitly assume the inclusion of spare parts.

6.4 Treating *value* as an Abbreviation

It will be possible to define *value* as an abbreviation in a modal context language which contains the *ist*. (Such reductions might be problematic in non-modal approaches, such as Attardi and Simi 1998.)

We first deal with the case where all contexts have the same domains. We define *value* as an abbreviation:

(60) $\quad value(c, x) = y \quad \equiv \quad (\forall z)y = z \equiv ist(c, x = z)$

Eliminating the *value* abbreviation, the above formulas are equivalent to:

(61) $\quad c_{ps} : \quad (\forall xy)ist(c_{GE}, y = price(x)) \equiv y = GE\text{-}price(x)$

(62) $\quad c_{ps} : \quad (\forall xy)ist(c_{NAVY}, y = price(x)) \equiv y = GE\text{-}price(x) +$
$$GE\text{-}price(spares(c_{NAVY}, x))$$

(63) $\quad c_{ps} : \quad (\forall xy)ist(c_{AF}, y = price(x)) \equiv$
$$y = f(x, GE\text{-}price(x), GE\text{-}price(spares(c_{AF}, x)))$$

6.5 Existence as a Predicate

Not all objects will typically exist in all contexts. To deal with this phenomena, we introduce existence as a predicate, $E(c, x)$. Intuitively,

$$E(c, x)$$

is true iff the object denoted by x exists in the context denoted by c.

If the domains of all the contexts are not the same, then formulas 60—63 are not intuitively correct. Instead, a domain precondition needs

to be added to all the formulas. For example instead of formula 61, we would write

(64) $c_{\mathrm{ps}}:\quad (\forall xy)(E(c,x) \wedge E(c,y)) \supset$

$$(ist(c_{\mathrm{GE}}, y = price(x)) \equiv y = \textit{GE-price}(x))$$

The main implication connective in this formula might not be classical (see Buvač et al. 1995).

Note however, if we simply change the abbreviation of *value* to

(65) $value(c,x) = y\quad \equiv\quad E(c,x) \supset (\forall z)(E(c,z) \supset$

$$(y = z \equiv ist(c, x = z)))$$

then the axioms involving *value* (axioms 40-42 and 58-59) will still be true. In other words, the previous domain formalizations remain unaltered. To verify this, note that substituting this new definition for value (given in formula 65) into formula 40 gives us formula 64, rather then formula 61.

We also need to assert that the problem solving context c_{ps} contains all the objects present in the other contexts which are involved in the particular problem solving process. In some outer context $c0$ we would write:

(66) $c0:\quad (\forall c)\textit{involved-in-ps}(c) \supset (\forall x)(E(c,x) \supset E(c_{\mathrm{ps}}, x)).$

In both cases mentioned above, the rules of entering and exiting a context for the *value* function will follow from the general rules enter and exit for the *ist*.

This approach is similar to the treatment of existence in quantificational modal logics which are based on free logics; see Thomason 1970. Many of the philosophical issues involved in quantifying across worlds with different domains apply also here; see Kripke 1971 for discussion. For example, it is well known from quantificational modal logic that the classical rule of universal instantiation is not valid for non-rigid terms. Logics of context which allow functions, will thus need to have a restricted form of universal instantiation.

7 Combining Planning Contexts

Integrating plans, which were not originally designed to be used together, is a task that frequently comes up in real world applications. However, this task is typically performed by humans. Contexts enable us to formalize this style of reasoning, thus providing a logical basis for developing computer programs which will be able to mechanically integrate plans produced by different systems.

The main contribution of this section is showing that the techniques

we have used for combining data bases or knowledge bases (§4–§6) can be directly applied in combining of plans produced by different planners using different languages.

See any AI textbook, say Rich and Knight 1991, Russell and Norvig 1995 or the forthcoming Nilsson 1998, for a description of the planning problem.

7.1   Combining Planning Languages

The basic idea is to represent each subplan as though it was developed in a context. These contexts will differ, either slightly or greatly. The terminology within each subcontext is likely to be specialized, and making the plans work together requires some generalization. For this purpose we use lifting formulas, i.e. formulas that relate the propositions and terms in subcontexts to propositions and terms in an outer context.

Here is an example. Assume that a route planner, like the route optimization program of the TRAINS project Allen et al. 1995, and a supply planner, like the transportation scheduler developed at Kestrel Smith et al. 1995, have been developed independently by different groups. Given a source and a destination, the route planner will find the best route between these places. It however, has no notion of which supplies need to be transported and no notion of time. The supply planner keeps track of the supplies of some economic system and informs us which supplies need to be moved at any given time. We assume that the supply planner has no knowledge about the routes that the supplies need to travel to reach their destination.

To fill in a work order we need to integrate the information produced by the supply planner with that of the route planner. Assume that the supply planner produces

(67)   *supply_planner* :   *transport*(*equipment1*,

   Rome, 11/6/95, *Frankfurt*, 1/20/96)

informing us that equipment1 needs to be transported from Rome Air Force base in New York on 11/6/95, to Frankfurt on 1/20/96. The context constant *supply_planner* denotes the context in which supply planner operates and reports its results. Now assume that the route planner tells us that the best route from Rome to Frankfurt is via New York City (NYC). This is represented in the context of the route planner, *route_planner*, by stating

(68)   *route_planner* :   [*Rome*, *NYC*, *Frankfurt*] =

   route(*Rome*, *Frankfurt*).

Note that *route* is a function returning a list which encodes the best

route. Integrating this information inside the problem solving context, *ps*, we get

(69) *ps* : *transport(equipment1, [Rome, NYC, Frankfurt],*

$\qquad$ *11/6/95, 1/20/96)*

stating that equipment1 needs to be transported by the route Rome–NYC–Frankfurt departing on 11/6/95 and arriving on 1/20/96. This information can now be entered into the work order. Note that the same predicate symbol, *transport*, is used in different ways in two different contexts: its arity and its arguments are different in the *supply_planner* context and in the *ps* context.

$\quad$ The context formalism enables us to capture this style of reasoning in logic. We write *lifting axioms* which describe how the information from different contexts can be integrated. In the above example the lifting formula is

(70) *ps* : $(\forall x)(\forall l1)(\forall l2)(\forall d1)(\forall d2)$

$\qquad$ *ist(supply_planner, transport(x, l1, d1, l2, d2))* $\supset$

$\qquad$ *transport(x, value(route_planner, route(l1, l2)), d1, d2).*

If the formula *transport(x, l1, d1, l2, d2)* is true in the context of the supply planner, then the formula

$\quad$ *transport(x, value(route_planner, route(l1, l2)), d1, d2)*

holds in the problem solving context. Intuitively, formula 70 expresses that if the supply planner states that some items x need to be transported leaving $l1$ on $d1$ and arriving to $l2$ on $d2$, and if

$\quad$ *value(route_planner, route(l1, l2))*

is reported by the route planner as the best route from $l1$ to $l2$, then the information which can be entered into the work order is

$\quad$ *transport(x, value(route_planner, route(l1, l2)), d1, d2).*

In other words, formula 70 specifies the integrating of a plan which involves the notions of time and supplies produced by the supply planner with the details involving a route produced the route planner. The lifting axiom 70 allows us to derive the plan given by formula 69 in the problem solving context *ps* from the plans given by formulas 67 and 68 in the contexts of their corresponding planners.

$\quad$ A term with a definite meaning in one context often needs translation when used in another context. Thus *Rome* may mean Rome NY in a data base of US Air Force bases but needs translation when a formula is lifted to a context of worldwide geography. Lifting formulas similar to 70 can be used to do this type of translation.

7.2 Discharging Kindness Assumptions

Any plan produced from the lifting axiom 70 makes numerous assumptions. For example, it assumes that the shortest path will always get the cargo to its destination on time. Although this assumption is usually valid, we can imagine a scenario in which an urgent delivery will need to take a longer route in order to get to its destination on time. We thus need to consider the timeliness of a path in scenarios which involve urgent deliveries.

In robotics, assumptions of this sort are commonly called *kindness assumptions*, cf. Nourbakhsh and Genesereth 1996, because they amount to assuming that the world is kind, i.e. that things will turn out in our favor most of the time. Kindness assumptions are a useful tool and are commonly made both when constructing and integrating plans. They allow us to focus on the aspects of a plan that seem to be relevant to the problem at hand and to disregard details which we assume will hold for that particular problem class. However, whenever kindness assumptions are made it is important to have a mechanism which enables us to discharge such assumptions and reason about their validity in cases when it is unclear whether they hold. The context formalisms enables us to do this in the framework of logic.

Assume that after deriving the plan in formula 69 (by integrating the plans from the route and supply planning contexts) we realize that the delivery is needed urgently. At this point our goal is to discharge the timeliness assumption and take the proposed path through NYC only if it gets equipment1 to Frankfurt on time. The desired plan, which is given in an urgent problem solving context ps_urgent, is thus

$$(71) \quad ps\_urgent: \quad timely\_route(11/6/95, 1/20/96,$$
$$[Rome, NYC, Frankfurt])) \supset transport(equipment1,$$
$$[Rome, NYC, Frankfurt], 11/6/95, 1/20/96)$$

where deciding whether *timely_route* holds will involve looking up airplane schedules and local delivery facilities in some data base. We are assuming that conditional plans, like formula 71, can be represented by the system. In the general case, formula 71 follows from formula 69 and the lifting axiom

$$(72) \quad ps\_urgent: \quad (\forall x)(\forall r)(\forall d1)(\forall d2) ist(ps, transport(x, r, d1, d2)) \supset$$
$$timely\_route(d1, d2, r) \supset transport(x, r, d1, d2).$$

In some planning instances we will want to consider the timeliness issues at the very outset. We can avoid using the original problem solving context *ps* by inferring a lifting theorem which integrates a plan from

the route planner and a plan from the supply planner to directly produce a plan in *ps_urgent*

(73) *ps_urgent* : $(\forall x)(\forall r)(\forall l1)(\forall l2)(\forall d1)(\forall d2)$

$(ist(supply\_planner, transport(x, l1, l2, d1, d2)) \wedge$

$timely\_route(d1, d2, value(route\_planner, route(l1, l2)))) \supset$

$transport(x, value(route\_planner, route(l1, l2)), d1, d2)$

Formula 73 logically follows from the lifting axioms given in formula 70 and formula 72.

7.3 Combining Heterogeneous Objects

It is often necessary to combine objects that were not designed to work together. These objects may be databases of facts, programs, hardware or plans. Even if the objects were intended to be used together, maintenance of adherence to their specifications and ensuring the consistency of terminology through time is often difficult. Each object was developed in a context, and these contexts differ, either slightly or greatly. The terminology within each subcontext is likely to be specialized, and making them work together requires some generalization.

We have described an approach based on our formal theory of context.

8 Representing Discourse

In this section we illustrate context change by showing how our formalism can be used to represent the context of a conversation in which terms have particular meanings that they wouldn't have in the language in general. The analysis that follows is along the lines of van Benthem 1996.

We examine question/answer conversations which are simply sequence of questions and answers. In this simple model we allow two types of questions:

propositional questions are used to inquire whether a proposition is true or false; they require a yes or no answer. In the language we introduce a special proposition *yes* which is used to answer these questions.

qualitative questions are used to find the objects of which a formula holds; in the language we introduce a unary predicate *answer* which holds of these objects.

In order to know what is being communicated in a discourse, as well as reason about a discourse in general, we need a way of representing the

discourse. To do this in the framework of the formal theory of context, we identify a new class of contexts, the *discourse contexts*. Discourse contexts are not only characterized by the sentences which are true in them but also by the intended meaning of their predicates, which might vary from one discourse context to the next.

We represent a discourse with a sequence of discourse contexts, each of which in turn represents the discourse state after an utterance in the discourse. Our attention is focused only on discourses which are sequences of questions and replies: $[q^1, r^1, q^2, r^2, \ldots, q^n, r^n]$. Thus, we can represent such a discourse with a sequence of discourse contexts:

$$[c_d, query(c_d, q^1), reply(query(c_d, q^1), r^1), \ldots$$

$$\ldots, reply(query(reply(\cdots reply(query(c_d, q^1), r^1) \cdots, r^{n-1}), q^n), r^n)]$$

s.t. (i) c_d^0 is some discourse context in which the initial question (q^1) was asked; (ii) the function *query* takes a question ϕ and some discourse context c_d (representing the discourse state before the question ϕ) and returns the discourse context representing the discourse state after asking the question ϕ in c_d; (iii) the function *reply* takes a reply ϕ and some discourse context c_d (representing the discourse state before before replying ϕ) and returns the discourse context representing the discourse state after replying ϕ in c_d. In order to reason about the discourse we now only need the properties of the functions *query* and *reply*. These will be made precise in the next subsection. Similar representation of discourse in logic is often used by linguists; eg. Stalnaker 1998, Beaver 1997.

Since we are not concerned with solving the syntactic and semantic problems addressed by the natural language community, we are assuming the system is given the discourse utterances in the form of logical formulas. This assumption is in line with McCarthy 1990a; in McCarthy's terminology we would say that the discourse has been processed by both the parser and the understander to produce a logical theory. Note that the process of producing this theory is not precisely defined, and it is not completely clear how much common sense information is needed to generate it. It might turn out that producing such a theory requires the solution of the problem we had set out to solve. But for the time being let us take a positive perspective and assume the discourse theory is given; for further discussion of this point see Israel 1990, Kameyama 1994, Thomason 1997, Buvač 1997. Note that our simple model does not claim to capture all aspects of discourse interpretation. We have refrained from modeling some phenomena that have been studied by semanticists and computational linguists. In particular, Discourse Representation Theory, Kamp 1981, includes a third aspect of discourse

interpretation, namely discourse entities known as *reference markers*. Reference markers, each of which can be accessible at a given point in a discourse, are now viewed as an essential element of most theories of context that deal with anaphoric reference. Furthermore, we have ignored pragmatic aspects relevant to discourse analysis; see Grosz and Sidner 1986. These include resolving references by keeping track of which objects are salient in a discourse, and inferring the intentions of agents based on their speech acts.

8.1 The Logic of *query* and *reply*

In this section we give the properties of the functions *query* and *reply*, which are central for representing question/answer discourses. Since the *query* and *reply* functions are treated in the style of situation calculus, we do not need to change our basic theory of context, but simply give the axioms that formalize the two functions.

Intuitively, the *query* function will set up a context in which the reply to the question will be interpreted. For example, the context resulting in asking some proposition p will have the property that *yes* in that context will be interpreted as p. Thus *query* only changes the semantic state of the discourse context. The *reply* function will do a simple update of information: the formulas true in the context resulting in replying p in c_d will be exactly those formulas which are conditionally true on p in c_d. Thus the *reply* function only changes the epistemic state of the discourse context. We now make these notions more precise.

The following axioms characterize the functions *query* and *reply*.

interpretation axiom (propositional) if ϕ is a closed formula, then

$$ist(query(c, \phi), \phi \equiv yes)$$

frame axiom (propositional) if ϕ is a closed formula, and *yes* does not occur in ψ, then

$$ist(c, \psi) \supset ist(query(c, \phi), \psi)$$

interpretation axiom (qualitative) if x is the only variable occurring free in ϕ, then

$$ist(query(c, \phi(x)), \phi(x) \equiv answer(x))$$

frame axiom (qualitative) if x is the only variable occurring free in ϕ, and *answer* does not occur in ψ, then

$$ist(c, \psi) \supset ist(query(c, \phi(x)), \psi)$$

reply axiom $\qquad ist(reply(c, \phi), \psi) \equiv ist(c, \phi \supset \psi)$

The predicate $answer(x)$ is used to return an answer to a qualitative

question, similar to the way it is commonly used to return a witness in a resolution theorem prover; see Green 1969.

We proceed to illustrate the axioms and their use with an example.

8.2 Example: Air Force–GE Discourse

We examine the following hypothetical discourse taking place between the Air Force and General Electric:

1. AF: Will you bid on the engine for the FX22?
2. GE: Yes.
3. AF: What is your bid?
4. GE: $4M.
5. AF: Does that include spares?
6. GE: Yes.

We transcribe the above discourse in our logic as a sequence of discourse contexts, s.t.

$c1 = query(c, \textit{will-bid-on}(engine(\text{FX22})))$

$c2 = reply(c1, yes)$

$c3 = query(c2, price(engine(\text{FX22}), x))$

$c4 = reply(c3, answer(\$4M))$

$c5 = query(c4, price(x) \equiv ist(c_{kb}, \textit{price-including-spares}))$

$c6 = reply(c5, yes)$

where c is the initial discourse context. To simplify presentation, in this section we take *price* to be a predicate; in §4 we have illustrated how it can be treated as a function by using *value* instead of *ist*.

8.3 Deriving Properties of the Air Force–GE Discourse

We now show some properties of the discourse which can be derived with our logic.

8.3.1 First Question: Propositional Case

The discourse begins with a propositional question. We show how they modify the discourse state.

Theorem ($c2$): $ist(c2, \textit{will-bid-on}(engine(\text{FX22})))$

Proof ($c2$):   Instantiating the first axiom for the propositional questions, we get

$ist(query(c, \textit{will-bid-on}(engine(\text{FX22}))),$

$\quad\quad \textit{will-bid-on}(engine(\text{FX22})) \equiv yes)$

which, by definition of c_1, can be written as

$$ist(c_1, \textit{will-bid-on}(engine(\text{FX22})) \equiv \textit{yes})$$

Instantiating the axiom for reply we have

$$ist(reply(c_1, yes), \textit{will-bid-on}(engine(\text{FX22}))) \equiv$$
$$ist(c_1, yes \supset \textit{will-bid-on}(engine(\text{FX22})))$$

and it follows from the two lines above that

$$ist(reply(c_1, yes), \textit{will-bid-on}(engine(\text{FX22})))$$

which by definition of c_2 we can write as

$$ist(c_2, \textit{will-bid-on}(engine(\text{FX22})))$$

$\square_{\mathbf{c2}}$

8.3.2 Second Question: Qualitative Case

The reasoning for this qualitative question is similar to the propositional question.

Theorem (c_4): $ist(c_4, price(engine(\text{FX22}), \$4M))$

Proof (c_4): We begin with an instance of the first axiom for qualitative questions

$$ist(query(c_2, price(engine(\text{FX22}), x)),$$
$$price(engine(\text{FX22}), x) \equiv answer(x))$$

which, by definition of c_3, can be written as

$$ist(c_3, price(engine(\text{FX22}), x) \equiv answer(x))$$

Instantiating the axiom for reply we have

$$ist(reply(c_3, answer(\$4M)), price(engine(\text{FX22}), \$4M)) \equiv$$
$$ist(c_3, answer(\$4M) \supset price(engine(\text{FX22}), \$4M))$$

and it follows from the two lines above that

$$ist(reply(c_3, answer(\$4M)), price(engine(\text{FX22}), \$4M))$$

which by definition of c_4 we can write as

$$ist(c_4, price(engine(\text{FX22}), \$4M))$$

$\square_{\mathbf{c4}}$

Due to the frame axioms, the conclusion established in the first question

$$ist(c_2, \textit{will-bid-on}(engine(\text{FX22})))$$

also holds in context c_4.

Theorem (frame): $ist(c2, will\text{-}bid\text{-}on(engine(FX22)))$

Proof (frame):   We first instantiate the second axiom for qualitative questions to get

$ist(c2, will\text{-}bid\text{-}on(engine(FX22))) \supset$
$ist(query(c2, price(engine(FX22), x)), will\text{-}bid\text{-}on(engine(FX22)))$

The two lines above imply

$ist(query(c2, price(engine(FX22), x)), will\text{-}bid\text{-}on(engine(FX22)))$

which, by definition of $c3$, can be written as

$ist(c3, will\text{-}bid\text{-}on(engine(FX22)))$

Now we apply the following instance of the reply axiom

$ist(reply(c3, answer(\$4M)), will\text{-}bid\text{-}on(engine(FX22))) \equiv$
$ist(c3, answer(\$4M) \supset will\text{-}bid\text{-}on(engine(FX22)))$

to get

$ist(reply(c3, answer(\$4M)), will\text{-}bid\text{-}on(engine(FX22)))$

which, by definition of $c4$, can be written as

$ist(c4, will\text{-}bid\text{-}on(engine(FX22)))$

□**frame**

8.3.3   Third Question: Dealing with Ambiguity

We are assuming that the predicate *price* is ambiguous in the discourse contexts since it can be ambiguously interpreted as either the predicate *price-including-spares* or as the predicate *price-not-including-spares* in some knowledge base. In the third question the predicate is disambiguated for context $c6$. This will allow us to prove that the GE bid on the FX22 engine is \$4M including spare parts. Note that we will have to state the above in the kb context because the discourse contexts are not expressive enough to distinguish between the price including spares and the price excluding spares (which in fact was the source of ambiguity).

Theorem (kb): $ist(c_{kb}, price\text{-}including\text{-}spares(engine(FX22), \$4M))$

Proof (kb):   By reasoning similar to the first question, we can conclude

$ist(c6, price(x, y)) \equiv ist(c_{kb}, price\text{-}including\text{-}spares(x, y))$

From the frame axioms we get

$ist(c6, price(engine(FX22), \$4M))$

similarly to the frame derivation in the second question. Now the theorem follows from the above formulas. $\square_{\mathbf{kb}}$

See Buvač 1996b for more details.

9 Transcending Contexts

Human intelligence involves an ability that no-one has yet undertaken to put in computer programs—namely the ability to *transcend* the context of one's beliefs.

That objects fall would be expected to be as thoroughly built into human mental structure as any belief could be. Nevertheless, long before space travel became possible, the possibility of weightlessness was contemplated. It wasn't easy, and Jules Verne got it wrong when he thought that there would be a turn-over point on the way to the moon when the travelers, who had been experiencing a pull towards the earth would suddenly experience a pull towards the moon.

In fact, this ability is required for something less than full intelligence. We need it to be able to comprehend someone else's discovery even if we can't make the discovery ourselves. To use the terminology of McCarthy and Hayes 1969, it is needed for the *epistemological* part of intelligence, leaving aside the heuristic.

We want to regard the system as being at any time within an implicit outer context; we have used $c0$ in this paper. Thus a sentence p that the program believes without qualification is regarded as equivalent to $ist(c0, p)$, and the program can therefore infer $ist(c0, p)$ from p, thus *transcending* the context $c0$. Performing this operation again should give us a new outer context, call it c_{-1}. This process can be continued indefinitely. We might even consider continuing the process transfinitely, for example, in order to have sentences that refer to the process of successive transcendence. However, I have no present use for that.

However, if the only mechanism we had is the one described in the previous paragraph, transcendence would be pointless. The new sentences would just be more elaborate versions of the old. The point of transcendence arises when we want the transcending context to relax or change some assumptions of the old. For example, our language of adjacency of physical objects may implicitly assume a gravitational field, e.g. by having relations of *on* and *above*. We may not have encapsulated these relations in a context. One use of transcendence is to permit relaxing such implicit assumptions.

The formalism might be further extended to provide so that in c_{-1} the whole set of sentences true in c_0 is an object $truths(c0)$.

Transcendence in this formalism is an approach to formalizing something that is done in science and philosophy whenever it is necessary to go from a language that makes certain assumptions to one that does not. It also provides a way of formalizing some of the human ability to make assertions about one's own thoughts.

The usefulness of transcendence will depend on there being a suitable collection of nonmonotonic rules for *lifting* sentences to the higher level contexts.

As long as we stay within a fixed outer context, it seems that our logic could remain ordinary first order logic. Transcending the outermost context seems to require a changed logic with what Tarski and Montague call *reflexion principles*. They use them for sentences like $true(p*) \equiv p$, e.g " 'Snow is white.' is true if and only if snow is white."

The above discussion concerns the epistemology of transcending contexts. The heuristics of transcendence, i.e. when a system should transcend its outer context and how, is entirely an open subject.

10 Relative Decontextualization

Quine [1969] uses a notion of "eternal sentence", essentially one that doesn't depend on context. This seems a doubtful idea and perhaps incompatible with some of Quine's other ideas, because there isn't any language in which eternal sentences could be expressed that doesn't involve contexts of some sort. We want to modify Quine's idea into something we can use.

The usefulness of eternal sentences comes from the fact that ordinary speech or writing involves many contexts, some of which, like pronoun reference, are valid only for parts of sentences. Consider, "Yes, John McCarthy is at Stanford University, but he's not at Stanford today". The phrase "at Stanford" is used in two senses in the same sentence. If the information is to be put (say) in a book to be read years later by people who don't know McCarthy or Stanford, then the information has to be decontextualized to the extent of replacing some of the phrases by less contextual ones.

The way we propose to do the work of "eternal sentences" is called *relative decontextualization*. The idea is that when several contexts occur in a discussion, there is a common context above all of them into which all terms and predicates can be lifted. Sentences in this context are "relatively eternal", but more thinking or adaptation to people or programs with different presuppositions may result in this context being transcended.

11   Mental States as Outer Contexts

A person's state of mind cannot be adequately regarded as the set of propositions that he believes—at least not if we regard the propositions as sentences that he would give as answers to questions. For example, as we write this we believe that George Bush is the President of the United States, and if we were entering information in a database, we might write

$$president(U.S.A.) = George.Bush.$$

However, my state of mind includes, besides the assertion itself, my reasons for believing it, e.g. he has been referred to as President in today's news, and we regard his death or incapacitation in such a short interval as improbable. The idea of a TMS (see Doyle 1979) or reason maintenance system is to keep track of the pedigrees of all the sentences in the database and keep this information in an auxiliary database, usually not in the form of sentences.

Our proposal is to use a database consisting entirely of *outer* sentences where the pedigree of an *inner* sentence is an auxiliary parameter of a kind of modal operator surrounding the sentence. Thus we might have the outer sentence

$$believe(president(U.S.A.) = George.Bush, because\ldots),$$

where the dots represent the reasons for believing that Bush is President.

The use of formalized contexts provides a convenient way of realizing this idea. In an outer context, the sentence with reasons is asserted. However, once the system has committed itself to reasoning with the proposition that Bush is President, it enters an inner context with the simpler assertion

$$president(U.S.A.) = George.Bush.$$

If the system then uses the assertion that Bush is President to reach a further conclusion, then when it leaves the inner context, this conclusion needs to acquire a suitable pedigree.

Consider a belief revision system that revises a database of beliefs solely as a function of the new belief being introduced and the old beliefs in the system. Such systems seem inadequate even to take into account the information used by TMS's to revise beliefs. However, it might turn out that such a system used on the outer beliefs might be adequate, because the consequent revision of inner beliefs would take reasons into account.

12 Remarks

1. Guha has put contexts into Cyc, largely in the form of microtheories. The *above-theory* example is a microtheory. See Guha 1991 for some of the details.

2. We have mentioned various ways of getting new contexts from old ones: by specializing the time or place, by specializing the situation, by making abbreviations, by specializing the subject matter (e.g. to U.S. legal history), by making assumptions and by specializing to the context of a conversation. These are all specializations of one kind or another. Getting a new context by transcending an old context, e.g. by dropping the assumption of a gravitational field, gives rise to a whole new class of ways of getting new contexts.

These are too many ways of getting new contexts to be treated separately.

3. We have used natural language examples in this article, although natural language is not our main concern. Nevertheless, we hope that formalizing context in the ways we propose may be useful in studying the semantics of natural language. Natural language exhibits the striking phenomenon that context may vary on a very small scale; several contexts may occur in a single sentence.

Consider the context of an operation in which the surgeon says, "Scalpel". In context, this may be equivalent to the sentence, "Please give me the number 3 scalpel".

4. $ist(c, p)$ can be considered a modal operator dependent on c applied to p. In this sense much of our analysis amounts to reasoning in a certain systems of modal logic or temporal logic; see Chellas 1980, Gabbay et al. 1993, Blackburn and de Rijke 1997.

In the propositional case, given a context language containing a set of contexts C, we can define a modal language containing modalities $\Box_1, \Box_2, \ldots$, one for each context from C. By replacing each occurrence of $ist(c_\beta, \phi)$ with $\Box_\beta \phi$, we can define a bijective translation function which to each formula of the propositional context logic assigns a well-formed modal formula. Based on this translation, Buvač et al. 1995 shows a reduction of the propositional logic of context to a propositional multi-modal logic. Similar results are obtained using proof theoretic tools in Giunchiglia and Serafini 1994.

However, these results do not carry over to the quantificational case. The quantificational logic of context, for example, enables us to state

that the formula ϕ is true in contexts which satisfy some property $p(x)$ as follows:$(\forall v)p(v) \supset ist(v, \phi)$. This formula has no obvious translation into standard multi-modal logic, and the meaning of such formulas which quantify over modalities is beyond the analysis commonly done in quantificational modal logic. See Buvač 1996a for details.

5. Proof theoretic approach to context has been emphasized by Richard W. Weyhrauch and Fausto Giunchiglia and his group. See Weyhrauch 1980, Weyhrauch et al. 1998, Giunchiglia 1993, Giunchiglia and Serafini 1994, Cimatti and Serafini 1995.

6. It would be useful to have a formal theory of the natural phenomenon of context, e.g.in human life, as distinct from inventing a form of context useful for AI systems using logic for representation. This is likely to be an *approximate theory* in the sense described in McCarthy 1979a. That is, the term "context" will appear in useful axioms and other sentences but will not have a definition involving "if and only if". Hayes 1997 outlines one such taxonomy of contexts.

7. Useful nonmonotonic rules for lifting will surely be more complex than the examples given. See Etherington and Crawford 1996 for context limited consistency check.

8. Theories along the lines of Barwise and Perry 1983 are in many ways similar to formal theories of context; van Benthem 1997 gives one comparison. Menzel 1996, Akman and Surav 1997 represent context using the tools of situation theory.

9. Gabbay 1996 proposes fibred semantics as a way of "weaving of logics". For a comparison of this approach to the formal theories of context, see Gabbay and Nossum 1997.

13 Conclusion

Our main motivation for formalizing contexts is to deal with the problem of generality in AI. We want to be able to make AI systems which are never permanently stuck with the concepts they use at a given time because they can always transcend the context they are in. Such a capability would allow the designer of a reasoning system to include only such phenomena as are required for the system's immediate purpose, while retaining the assurance that if a broader system is required later, "lifting axioms" can be devised to restate the facts from the narrower context to the broader one, with qualifications added as necessary. Thus,

a necessary step in the direction of addressing the problem of generality in AI is providing a language which enables representing and reasoning with multiple contexts and expressing lifting axioms. In this paper we provide such a language.

The goal is that no matter what corners the specialists paint themselves into, what they do can be *lifted* out and used in a more general context.

For an overview of the AI research on formalizing context see Akman and Surav 1996. For technical papers on context in AI and Linguistics see the following special issues of journals: Perlis 1995, Iwanska and Zadrozny 1997, Buvač and Kameyama 1998.

References

Akman, Varol, and Mehmet Surav. 1996. Steps toward Formalizing Context. *AI Magazine* 17(3):55–72.

Akman, Varol, and Mehmet Surav. 1997. The use of situation theory in context modeling. *Computational Intelligence* 13(3). To appear.

Allen, James F., Lenhart K. Schubert, George Ferguson, Peter Heeman, Chung Hee Hwang, Tsuneaki Kato, Marc Light, Nathaniel G. Martin, Bradford W. Miller, Massimo Poesio, and David R. Traum. 1995. A Formalization of Viewpoints. *Journal of Experimental and Theoretical AI* 7:7–48.

Amati, Gianni, and Fiora Pirri. 1997. Contexts as relativized definitions: a formalization via fixed points. In *Context in KR & NL*. MIT, Cambridge, Mass. AAAI Fall Symposium. To appear.

Attardi, Giuseppe, and Maria Simi. 1998. Communication across Viewpoints. *Journal of Logic Language and Information.* To appear.

Barwise, Jon, and John Perry. 1983. *Situations and Attitudes.* Cambridge: MIT Press.

Beaver, David Ian. 1997. Presupposition. In *Handbook of Logic and Language*, ed. Johan van Benthem and Alice ter Meulen. Elsevier—MIT Press.

van Benthem, Johan. 1996. *Exploring Logical Dynamics.* CSLI Publications. Distributed by Cambridge University Press.

van Benthem, Johan. 1997. Changing Contexts and Shifting Assertions. In *this volume.*

Blackburn, Patrick, and Maarten de Rijke. 1997. *Specifying Syntactic Structures.* CSLI Publications. Distributed by Cambridge University Press.

Buvač, Saša. 1996a. Quantificational Logic of Context. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence.*

Buvač, Saša. 1996b. Resolving lexical ambiguity using a formal theory of context. In *Semantic Ambiguity and Underspecification*, ed. Kees Van Deemter and Stanley Peters. CSLI Publications. Distributed by Cambridge University Press: CSLI Lecture Notes Number 55.

Buvač, Saša. 1997. Pragmatical Considerations on Logical AI. In *Context in KR & NL*. MIT, Cambridge, Mass. AAAI Fall Symposium. To appear.

Buvač, Saša, Vanja Buvač, and Ian A. Mason. 1995. Metamathematics of Contexts. *Fundamenta Informaticae* 23(3).

Buvač, Saša, and Richard Fikes. 1995. A Declarative Formalization of Knowledge Translation. In *Proceedings of the ACM CIKM: The 4th International Conference on Information and Knowledge Management*.

Buvač, Saša, and Megumi Kameyama. 1998. Introduction: Toward a Unified Theory of Context? *Journal of Logic Language and Information*. Special issue on Context in Linguistics and AI. To appear.

Buvač, Saša, and John McCarthy. 1996. Combining Planning Contexts. In *Advanced Planning Technology—Technological Achievements of the ARPA/Rome Laboratory Planning Initiative*, ed. Austin Tate. AAAI Press.

Chellas, Brian F. 1980. *Modal Logic, an Introduction*. Cambridge University Press.

Cimatti, A., and L. Serafini. 1995. Multi-Agent Reasoning with Belief Contexts: the Approach and a Case Study. In *Intelligent Agents: Proceedings of 1994 Workshop on Agent Theories, Architectures, and Languages*, ed. M. Wooldridge and N. R. Jennings, 71–85. Lecture Notes in Computer Science, No. 890. Springer Verlag. Also IRST-Technical Report 9312-01, IRST, Trento, Italy.

Collet, C., M. Huhns, and W. Shen. 1991. Resource Integration Using a Large Knowledge Base in Carnot. *IEEE Computer* 24(12):55–62.

Doyle, Jon. 1979. A Truth Maintenance System. *Artificial Intelligence* 12:231–272.

Etherington, David W., and James M. Crawford. 1996. Toward Efficient Default Reasoning. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence*.

Farquhar, Adam, Angela Dappert, Richard Fikes, and Wanda Pratt. 1995. Integrating Information Sources using Context Logic. Technical Report KSL-95-12. Knowledge Systems Laboratory, Stanford University. Also appears in the 1995 AAAI Spring Symposium on Information Gathering from Distributed, Heterogeneous Environments.

Gabbay, D. M., C. J. Hogger, and J. A. Robinson (ed.). 1993. *Handbook of Logic in Artificial Intelligence and Logic Programming, Volume 1 and Volume 4*. Oxford University Press.

Gabbay, Dov. 1996. Fibred semantics and the weaving of logics: Part I: Modal and intuitionistic logics. *Journal of Symbolic Logic* 61(4):1057–1120.

Gabbay, Dov, and Rolf Nossum. 1997. Structured Contexts with Fibred Semantics. In *Proceedings of the International and Interdisciplinary Conference on Context*.

Giunchiglia, F. 1993. Contextual reasoning. *Epistemologia, special issue on I Linguaggi e le Macchine* XVI:345–364. Short version in Proceedings IJ-

CAI'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., and L. Serafini. 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.

Green, C. Cordell. 1969. Theorem Proving by Resolution As a Basis for Question Answering Systems. In *Machine Intelligence 4*, ed. Bernard Meltzer and Donald Michie. 183–205. Edinburgh, Scotland: Edinburgh University Press.

Grosz, Barbara J., and Candace L. Sidner. 1986. Attention, Intention, and the Structure of Discourse. *Computational Linguistics* 12:175–204.

Guha, R. V. 1991. *Contexts: A Formalization and Some Applications*. Doctoral dissertation, Stanford University. Also technical report STAN-CS-91-1399-Thesis, and MCC Technical Report Number ACT-CYC-423-91.

Hayes, Pat. 1997. The Contextual Ambiguity of "Context". In *Context in KR & NL*. MIT, Cambridge, Mass. AAAI Fall Symposium. To appear.

Israel, D. J. 1990. On Formal versus Commonsense Semantics. In *Theoretical Issues in Natural Language Processing*, ed. Y. Wilks. Hillsdale, NJ: Lawrence Erlbaum Associates.

Iwanska, Lucja, and Wlodek Zadrozny. 1997. Special issue on Context in Natural Language Processing. *Computational Intelligence* 13(3). To appear.

Kameyama, Megumi. 1994. Indefeasible semantics and defeasible pragmatics. In *Quantifiers, Deduction, and Context*, ed. Makoto Kanazawa, Christopher Pinon, and Henriette de Swart. CSLI, Stanford, CA. Also SRI International Technical Note 544 and CWI Report CS-R9441,1994.

Kamp, Hans. 1981. A Theory of Truth and Discourse Representation. In *Formal Methods in the Study of Language*, ed. J. Groenendijk, T. Janssen, and M. Stokhof. Amsterdam: Mathematical Centre Tracts 135.

Ketonen, Jussi, and Joseph S. Weening. 1984. EKL—An Interactive Proof Checker: User's Reference Manual. Technical report. Computer Science Department, Stanford University, Stanford, California.

Kripke, Saul A. 1971. Semantical Considerations on Modal Logic. In *Reference and Modality*, ed. Leonard Linsky. Oxford University Press. Oxford Readings in Philosophy.

McCarthy, John. 1979a. Ascribing mental qualities to machines. In *Philosophical Perspectives in Artificial Intelligence*, ed. Martin Ringle. Humanities Press. Reprinted in McCarthy 1990b.

McCarthy, John. 1979b. First Order Theories of Individual Concepts and Propositions. In *Machine Intelligence*, ed. Donald Michie. Edinburgh: Edinburgh University Press. Reprinted in McCarthy 1990b.

McCarthy, John. 1987. Generality in Artificial Intelligence. *Comm. of ACM* 30(12):1030–1035. Also in *ACM Turing Award Lectures, The First Twenty Years*, ACM Press, 1987; and reprinted in McCarthy 1990b.

McCarthy, John. 1990a. An Example for Natural Language Understanding and the AI Problems it Raises. In *Formalizing Common Sense: Papers by John McCarthy*. 355 Chestnut Street, Norwood, NJ 07648: Ablex Publishing Corporation.

McCarthy, John. 1990b. *Formalizing Common Sense: Papers by John McCarthy*. 355 Chestnut Street, Norwood, NJ 07648: Ablex Publishing Corporation.

McCarthy, John. 1993. Notes on Formalizing Context. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*.

McCarthy, John, and Patrick Hayes. 1969. Some philosophical problems from the standpoint of artificial intelligence. In *Machine Intelligence*, ed. B. Meltzer and D. Michie. 463–502. Edinburgh: Edinburgh University Press. Reprinted in McCarthy 1990b.

Menzel, Christopher. 1996. Contexts and Information. In *The Third Symposium on Logical Formalizations of Commonsense Reasoning*. Stanford University. Preliminary report; updated versions available at http://philebus.tamu.edu/~cmenzel.

Nayak, P. Pandurang. 1994. Representing Multiple Theories. In *Proceedings of the Twelfth National Conference on Artificial Intelligence*.

Nilsson, Nils J. 1998. *Artificial Intelligence: A New Synthesis*. Morgan Kaufmann. To appear.

Nourbakhsh, Illah R., and Michael R. Genesereth. 1996. Making Assumptions to Simplify Planning. *Autonomous Robots* 1(3).

Perlis, D. 1995. Papers on Context: Theory and Practice. *Fundamenta Informaticae* 23(3).

Prawitz, Dag. 1965. *Natural Deduction: A Proof-Theoretical Study*. Stockholm, Göteborg, Uppsala: Almqvist & Wiksell Boktryckeri AB.

Rich, Elaine, and Kevin Knight. 1991. *Artificial Intelligence*. McGraw-Hill, Inc. 2 edition.

Russell, Stuart, and Peter Norvig. 1995. *Artificial Intelligence, a Modern Approach*. Prentice Hall.

Shoham, Yoav. 1991. Varieties of Context. In *Artificial Intelligence and Mathematical Theory of Computation: Papers in Honor of John McCarthy*, ed. Vladimir Lifschitz. Academic Press.

Smith, D. R., E. A. Parra, and S. J. Westfold. 1995. Synthesis of High-Performance Transportation Schedulers. Technical Report KES.U.95.1. Kestrel Institute. http://kestrel.edu/www/publications.html.

Stalnaker, Robert. 1998. On the Representation of Context. *Journal of Logic Language and Information*. To appear.

Thomason, Richmond H. 1970. Some Completeness Results for Modal Predicate Calculi. In *Philosophical Problems in Logic*, ed. Karel Lambert. D. Reidel Publishing Company/ Dordrecht-Holland. Synthese Library, Humanities Press, New York.

Thomason, Richmond H. 1997. Nonmonotonicity in Linguistics. In *Handbook of Logic and Language*, ed. Johan van Benthem and Alice ter Meulen. Elsevier—MIT Press.

Troelstra, A. S., and H. Schwichtenberg. 1996. *Basic Proof Theory*. Cambridge University Press. Cambridge Tracts in Theoretical Computer Science 43.

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

Weyhrauch, Richard W., Marco Cadoli, and Carolyn L. Talcott. 1998. Using Abstract Resources to Control Reasoning. *Journal of Logic Language and Information*. To appear.