NOTES ON FORMALIZING CONTEXT *

John McCarthy
Computer Science Department
Stanford University
Stanford, CA 94305
U.S.A.
(email: jmc@cs.stanford.edu)

Abstract

These notes discuss formalizing contexts as first class objects. The basic relation is $ist(c, p)$. It asserts that the *proposition* p is true in the *context* c. The most important formulas relate the propositions true in different contexts. 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. Various notations are considered, but these notes are tentative in not proposing a single language with all the desired capabilities.

1 Introduction

These notes contain some of the reasoning behind the proposals of [McCarthy, 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 I will feel free to use formulas in discussions of different applications that aren't always compatible with each other. [While I dithered, R.V. Guha wrote his dissertation.]

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 me 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.

*This work was partly supported by DARPA contract NAG2-703.

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 mattters 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

$$c' : \qquad ist(c, p), \tag{1}$$

which are to be taken as assertions that the proposition p is true in the context c, itself asserted in an outer context c'. (I have adopted Guha's [Guha, 1991] notation rather than that of [McCarthy, 1987], because he built his into Cyc, and it was easy for me to change mine. 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.

$$\begin{aligned} c0: \quad & ist(context\text{-}of(\text{“Sherlock Holmes stories”}), \\ & \text{“Holmes is a detective”}) \end{aligned}$$

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

$$c0: \quad ist(context\text{-}of(\text{"U.S. legal history"}), \text{"Holmes is a Supreme Court Justice"}).$$

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 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 assertions 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:

$$c0: \quad value(context\text{-}of(\text{"Sherlock Holmes stories"}), \text{"number of Holmes's wives"}) = 0$$

whereas

$$c0: \quad value(context\text{-}of(\text{"U.S. legal history"}), \text{"number of Holmes's wives"}) = 1.$$

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.

2 Relations among Contexts

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

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

$$
\begin{aligned}
c0: \quad & ist(specialize\text{-}time(t,c), at(jmc, Stanford)) \\
& \equiv ist(c, at\text{-}time(t, at(jmc, Stanford))).
\end{aligned}
$$

Here $at\text{-}time(t,p)$ is the assertion that the proposition p holds at time t. We call this a *lifting* relation. It is convenient to write $at\text{-}time(t, foo(x,y,z))$ rather than $foo(x,y,z,t)$, 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. $slowly(on\text{-}foot(go))$ rather than $go(on\text{-}foot, slowly)$.

Instead of using the function $specialize\text{-}time$, it may be convenient to use a predicate $specializes\text{-}time$ and an axiom

$$
\begin{aligned}
c0: \quad & specializes\text{-}time(t, c1, c2) \wedge ist(p, c1) \\
& \supset ist(c2, at\text{-}time(t,p)).
\end{aligned}
$$

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 of section 3.

3. There is a general relation $specializes$ between contexts. We say $specializes(c1,c2)$ when $c2$ involves no more assumptions than $c1$ and every *proposition* meaningful in $c1$ is translatable into one meaningful in $c2$. We have nonmonotonic relations

$$specializes(c1,c2) \wedge \neg ab1(p, c1, c2) \wedge ist(c1,p) \supset ist(c2,p).$$

and

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

This gives nonmonotonic inheritance of ist in both from the subcontext to the supercontext and vice versa. More useful is the case when the sentences must change when lifted. See below 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 $wire117$ 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 Leaving Contexts

Suppose we have the sentence $ist(c, p)$. We can then *enter* the context c and *infer* the sentence p. 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.

Another application of entering a context has to do with quantifiers. It involves a distinguished predicate $present(c, exp)$, where exp names an object. If we have

$$\forall x(present(c, x) \supset P(x)),$$

then when we enter c, then a special inference rule associated with the predicate *present* gives

$$\forall x P(x).$$

Likewise if we have shown

$$\exists x P(x)$$

within the context c, we can infer

$$\exists x(present(c, x) \wedge P(x)).$$

We could get similar effects by associating a domain (call it $domain(c)$) with each context c.

I'm 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. We plan to fix this up in some way later, either by introducing terms like $Ist(c, p)$ or by using a modified logic. Actually the same problem arises for p itself; the occurrence of p in $ist(c, p)$ might have to be syntactically distinct from the occurence of p standing by itself.

4 Rules for Lifting

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

$above\text{-}theory:$

$$(\forall xy)(on(x, y) \supset above(x, y)) \tag{2}$$

$$(\forall xyz)(above(x, y) \wedge above(y, z) \supset above(x, z)) \tag{3}$$

etc.

which stands for

$$c0: \qquad ist(above\text{-}theory, (\forall xy)(on(x, y) \supset above(x, y))) \tag{4}$$

etc.

We want to apply *above-theory* in a context c in which *on* and *above* have a third argument denoting a situation. In the following formulas, we put the context in which the formula is true to the left followed by a colon. $c0$ denotes an outer context in which formulas not otherwise qualified are true. The next section has more about $c0$. Suppose that in context c we have

$$c: \qquad (\forall xys)(on(x,y,s) \equiv ist(c1(s), on(x,y))), \tag{5}$$

and

$$c: \qquad (\forall xys)(above(x,y,s) \equiv ist(c1(s), above(x,y))), \tag{6}$$

etc.,

thus associating a context $c1(s)$ with each situation s. We also need

$$c0: \qquad ist(c, (\forall p\ s)(ist(above\text{-}theory, p) \supset ist(c1(s), p))), \tag{7}$$

which abbreviates to

$$c: \qquad (\forall p\ s)(ist(above\text{-}theory, p) \supset ist(c1(s), p)), \tag{8}$$

and asserts that the facts of *above-theory* all hold in the contexts associated with situations. Mike Genesereth points out that this necessarily involves quantifying into an *ist*. Now suppose we have the specific fact

$$c0: \qquad ist(c, on(A, B, S0)) \tag{9}$$

asserting that block A is on block B in a specific situation $S0$, and we want to derive $ist(c, above(A, B, S0))$. We proceed as follows.

First use (5) to get

$$c: \qquad ist(c1(S0), on(A, B)). \tag{10}$$

Now we enter $c1(S0)$ and get

$$c1(S0): \qquad on(A, B). \tag{11}$$

From (4) and (8) we conclude

$$c: \qquad ist(c1(S0), (\forall xy)(on(x,y) \supset above(x,y))), \tag{12}$$

from which entering $c1(S0)$ gives

$$c1(S0): \qquad (\forall xy)(on(x,y) \supset above(x,y)). \tag{13}$$

(11) and (13) give

$$c1(S0): \qquad above(A, B), \tag{14}$$

holding in context $c1(S0)$. We can now either continue reasoning in $c1(S0)$ or leave $c1(S0)$ and get

$$c: \qquad ist(c1(S0), above(A, B)) \tag{15}$$

and using (6)

$$c: \qquad above(A, B, S0) \tag{16}$$

and finally

$$c0: \qquad ist(c, above(A, B, S0)). \tag{17}$$

In this derivation we used a function giving a context $c1(s)$ depending 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 notation of [Buvac and Mason, 1993] or the ideas of [McCarthy, 1979b], but all have disadvantages. At present we are more attached to the derivation than to any specific logical system and consider preferable a system in which the above derivation is preserved with as little change as possible.

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 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 travellers, 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 asumptions 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.

6 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.

7 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 I write this I believe that George Bush is the President of the United States, and if I were entering information in a database, I might write

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

However, my state of mind includes, besides the asertion itself, my reasons for believing it, e.g. he has been referred to as President in today's news, and I regard his death or incapacitation in such a short interval as improbable. The idea of a TMS 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

$$president(U.S.A.) = GeorgeBush.$$

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.

8 Short Term Applications

We see the use of formalized contexts as one of the essential tools for reaching human level intelligence by logic based methods. However, we see formalized contexts as having shorter term applications.

- 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.

- Suppose the Air Force and General Electric Co. each have databases that include prices of jet engines and associated equipment. The items overlap in that jet engines that General Electric sells to the Air Force are included. Suppose further that the databases are not entirely compatible because the prices are based on different assumptions about spare parts and warranty conditions. Now suppose that the databases are to be used together by a program that must check whether the Air Force database is up-to-date on General Electric prices.

  Our idea is that corresponding to each database is a context, e.g. *context-GE-engine-prices* and *context-AF-engine-prices*. The program, however, must work with a context we may call *context-GE-AF-engine-prices*. Its language allows statements with auxiliary information about what is included in the price of an item. Suitable *lifting rules* allow translating the sentences of the two other databases into this more comprehensive context.

9 Remarks

1. 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.

2. We have used natural language examples in this article, although natural language is not our main concern. Nevertheless, I 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".

3. $ist(c, p)$ can be considered a modal operator dependent on c applied to p. This was explored in [Shoham, 1991].

4. 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".

5. Useful nonmonotonic rules for lifting will surely be more complex than the examples given.

Acknowledgments

The development of these ideas has benefitted from discussions with Saša Buvač, Tom Costello, Mike Genesereth, Fausto Giunchiglia and R. V. Guha. Guha wrote his thesis [Guha, 1991] while this article was going through many versions as the ideas developed, and the mutual influences cannot be specified.

References

[Buvac and Mason, 1993] **Saša Buvač and Ian A. Mason.**: "Propositional logic of context", in *Proceedings of the Eleventh National Conference on Artificial Intelligence*, 1993. To appear.

[Guha, 1991] **Guha, R. V.**: *Contexts: A Formalization and Some Applications*, Stanford PhD Thesis, 1991.

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

[McCarthy and Hayes, 1969] **McCarthy, John and P.J. Hayes**: "Some Philosophical Problems from the Standpoint of Artificial Intelligence", in D. Michie (ed), *Machine Intelligence 4*, American Elsevier, New York, NY, 1969. Reprinted in [McCarthy, 1990].

[McCarthy, 1979a] **McCarthy, John (1979a)**: "Ascribing Mental Qualities to Machines" in *Philosophical Perspectives in Artificial Intelligence*, Ringle, Martin (ed.), Harvester Press, July 1979. Reprinted in [McCarthy, 1990].

[McCarthy, 1979b] **McCarthy, John (1979b)**: "First Order Theories of Individual Concepts and Propositions", in Michie, Donald (ed.) *Machine Intelligence 9*, (University of Edinburgh Press, Edinburgh). Reprinted in [McCarthy, 1990]..

[McCarthy, 1987] **McCarthy, John (1987)**: "Generality in Artificial Intelligence", *Communications of the ACM*. Vol. 30, No. 12, pp. 1030-1035. Also in *ACM Turing Award Lectures, The First Twenty Years*, ACM Press, 1987. Reprinted in [McCarthy, 1990]..

[McCarthy, 1990] **McCarthy, John**: *Formalizing Common Sense*, Ablex, Norwood, New Jersey, 1990.

[Quine, 1969] **Quine, W. V. O.**: "Propositional Objects", in *Ontological Relativity and other Essays*, Columbia University Press, New York, 1969.

[Shoham, 1991] **Shoham, Y.**: "Varieties of Context", in *Artificial Intelligence and Mathematical Theories of Computation*, Academic Press, San Diego and London, 1991.

/@sail.stanford.edu:/u/ftp/pub/jmc/context-2.tex: created 1991 winter, latexed 1996 Mar 24 at 10:32 a.m.