Context as a Spurious Concept

Graeme Hirst

I take issue with formalizations of context in artificial intelligence, primarily the formalization by McCarthy and Buvač, that regard context as an undefined primitive whose formalization can be the same in many different kinds of AI tasks. In particular, any theory of context in natural language must take the special nature of natural language into account and cannot regard context simply as an undefined primitive. I show that there is no such thing as a coherent theory of context *simpliciter*—context pure and simple—and that context in natural language is not the same kind of thing as context in knowledge representation. In natural language, context is *constructed* by the speaker and the interpreter, and both have considerable discretion in so doing. Therefore, a formalization based on pre-defined contexts and pre-defined 'lifting axioms' cannot account for how context is used in real-world language.

1 CONTEXT IS AS CONTEXT DOES

> *The solution to any problem in AI may be found in the writings of Wittgenstein, though the details of the implementation are sometimes rather sketchy.*
>
> —*R.M. Duck-Lewis*

My purpose in this paper is to show that there is no such thing as a coherent theory of context *simpliciter*—context pure and simple; that context in natural language is not the same kind of thing as context in knowledge representation; and that any theory of context in natural language must take the special nature of natural language into account and cannot regard context simply as an undefined primitive.

The concept of 'context' has qualities rather like those of the concept of 'stressor'. A *stressor* is simply anything that causes stress—a deviation or distortion of a system from its normal state. A traffic jam might be a stressor upon one particular person, but have no such effect, or a different effect, upon another person. A sport fishery might be a stressor upon a boreal aquatic ecosystem. Thus, stressors are defined solely in terms of their effects—in fact, in terms of their effects on any particular person or system. What acts as a stressor for one person or system might have no effect at all on someone or something else. Just about anything can, in principle, be a stressor of something else. We can't, therefore, have any theory of the concept of 'stressor' *simpliciter*. There is no procedure by which we can determine whether or not a particular entity is a stressor just by looking at its properties or attributes. All we can do is actually apply the putative stressor to possible victims, and see if at least one of them experiences stress. Certainly, we can talk of 'frequent stressors' or 'likely stressors' of various kinds of stressees, and we can have theories of what kinds of objects these are and recognition procedures for them. But it's only when we particularize in just this way that we can have such theories and procedures.

In other words, 'stressor' is not a natural kind. And neither is 'context'. Like 'stressor', 'context' is a concept that is defined solely in terms of effects in a given situation.

Just about anything can, in principle, be a context. Whether something actually *is* a context can be determined only by its effect (which I'll describe a little while later). And what is a context in one particular case might not be a context in another case. What this means is that we can't have any theory of the nature of a context. There is no procedure by which we can determine whether or not a particular entity is a context just from looking at its properties or attributes. All we can do is apply the putative context to possible victims, and see if at least one of them experiences a contextual effect.

Consequently, any approach to 'context' *simpliciter* that tries to or purports to reify it, formalize it, or just speak of different views of it is inherently misguided.

McCarthy and Buvač [1997], for example, explicitly decline to give any definition of 'context'—it's as undefined as an element of a group, says McCarthy in another paper [1996]. But then they proceed to stipulate—despite the lack of definition—that contexts can be formalized as first-class objects, all of the same formal type, that they're things that propositions can be true in, and that they're things that can be entered and exited and nested. McCarthy and Buvač seem to see contexts as *containers* of some kind, at least metaphorically speaking. But that's just an assumption that they make, and it's an assumption that they make so deeply that they never even refer to it explicitly. My point in this paper is that 'context' simply doesn't permit this kind of approach.

Now, understand here that I am *not* opposing the general enterprise of formalizing abstracta. On the contrary, it's an enterprise that I engage in myself—for example in Hirst 1995, I wrote of "differences as first-class objects". Rather, I *am* suggesting that formalization, when appropriate, should be just about the final step, rather than the first step, in the understanding of a putative concept. And in the case of 'context', it's clear that we are still in a *very* early stage of understanding.

McCarthy [1996] suggests that a mathematical logic of contexts would be analogous to the mathematical theory of groups. But he himself points out that group theory arose from observations that the algebraic properties of integers under addition were the same as those of the rationals under multiplication, and the appropriate abstractions could then be made. In the case of context, however, we don't even have the observations yet. Sure, we can devise a nice new formal logic, and even call it a logic of context rather than, say, a logic of snibs or snecks or some gensymed word. But if the logic is to have something to say about what the English word *context* is about, then a little more work is in order. In his section entitled "Desiderata for a mathematical logic of context", even though he explicitly mentions applications in natural language, McCarthy lists just a few matters of formalization, and despite the heading, no desiderata deriving from the rather obvious need that a logic of context account for what context does in natural language, nor even the desideratum of finding out what such desiderata might be.

Although they side-step any definition of 'context', McCarthy and Buvač [1997], do say that contexts are "rich" objects. But they never open them up and look at their internal structure, preferring instead to follow a path analogous to the development of group theory; yet they want their work to be genuinely useful in AI applications. This is a bit like saying that a course on the algebra of groups, rings, and fields is the only qualification that anyone needs in order to become a professional accountant or bookkeeper.

2 INFORMAL NOTIONS OF CONTEXT

While many researchers in artificial intelligence talk about "context", or use representations that implicitly or explicitly act as "context" in some sense, the notion of context remains largely pre-theoretical. Indeed, even to speak of "pre-theoretical notions of context" implies that a theory of context *simpliciter* will eventuate in due course, and, as I've just said, I don't think that there can be any such thing. Instead, I'll speak of "informal" notions of context, and restrict the discussion in such a way as to make it possible to meaningfully theorize about contexts.

It doesn't go without saying that it's meaningful to speak at all of a unified notion of context in artificial intelligence. The word *context* really covers quite a board territory. In their excellent survey article on the formalization of context, Akman and Surav [1996] show that there are many different kinds and uses of context even just within artificial intelligence. Sure, there are similarities among these different kinds of context—that's why we use the same name for each—but it doesn't follow that everything we say about one kind will automatically be true of another kind. Maybe context in knowledge representation and context in natural language, for example, are two qualitatively different things. If so, then talking about 'context in knowledge representation and natural language' would be like talking about 'stressors of people and boreal aquatic ecosystems', and trying to find out just what a traffic jam has in common with a sport fishery. One of our jobs in research on context is to try to sort out the different meanings, and not just to presuppose that no such work is needed.

So I'll explain next why I believe that, indeed, context in knowledge representation and formal reasoning is not the same as context in natural language. It will follow from this that there can't be any useful theory or formalization of context *simpliciter*, because the behaviour of each kind of context is different.

But first, I need to make a distinction that I would have preferred to have made a little earlier. I want to distinguish between 'context' and 'element of context'. When I spoke earlier of "the effects of a context", what I really wanted to say was "the effects of one or more *elements* of context". But I couldn't actually say that earlier, because at that point I was still granting the idea of 'context' as a primitive. But now that we realize that we need to look at the internal structure of a context, we can talk about its individual elements and about their effects, either singly or in concert.

3 CONTEXT IN KNOWLEDGE REPRESENTATION

Considering first context in knowledge representation, I see no problem with McCarthy's basic idea [1987] of making axioms context-dependent in order to be able to state them at the most convenient or useful level of generality, nor with the suggestions as to the advantages that might be gained from doing this that are set out by Shoham [1991] and Akman and Surav [1996]. McCarthy and Buvač's well-known example [1997] in which different databases make different assumptions regarding the price of airplane components shows the benefits of this approach.

But at the same time it shows the limitations. The example involves formal, propositional reasoning, and the notion of a proposition being true or false in a context. It assumes that the assumptions made by the databases are static; and indeed, the exercise of writing *any* context-dependent axiom assumes that the 'home' context for the axiom is, in effect, pre-defined; that contexts can be usefully related by generalization and specialization; and that 'lifting axioms' can be predefined to relate truth in one context to truth in another. That's fine in a formal system, but it doesn't get us very far with language, to which I'll now turn.

4 CONTEXT IN NATURAL LANGUAGE

When it comes to talking about context in natural language, there is overwhelming consensus, I believe, on at least one point: context is a source of information that can be used (is used, should be used, may be used, must be used) by a language processor to reduce (or completely eliminate) ambiguity, vagueness, or underspecification in its interpretations of the utterances that it processes. That's one of the effects of context. It constrains interpretation.

In addition, context affects both what the speaker intends to say in the first place and how he or she goes about doing that. But there are really two kinds of context in this. One is the situation in which the speaker, as an agent, forms the intent to do something, which in this case is to communicate some message. This context constrains the agent's *intent*. The other is the context that the speaker uses as a source of information in creating that message, deciding exactly how it is to be expressed to the particular hearer. This context constrains the *form* of the communication and its exact *content*.

A point that immediately arises from this is that the use of context in natural language communication is a psychological construct that is not directly concerned with truth, but rather with interpretation and belief, with the *generation of meaning*. Propositional truth is involved only insofar as the interpreter may use their beliefs about what is and isn't true when they form an interpretation of some utterance. In natural language, a context is not something that propositions are "true in". It's something that interpretations are formed in, or, more precisely, formed *with*. An interpretation can be more or less vague or ambiguous, and more or less in accord with the speaker's intent, and if the interpretation is a proposition—which it needn't be, of course—then it might indeed have a truth value—though it needn't, of course.

A second point that arises is that the speaker has considerable discretion in the selection or construction of the context that is used in forming the utterance. That is, the speaker can (under constraints that I'll get to shortly) choose which potential elements of context to attend to and use and which ones to ignore. Likewise, the interpreter has considerable discretion in the construction of the interpretation—and, notwithstanding anything the speaker does, has discretion even in the selection or construction of the context that is used to create the interpretation. That is, the interpreter can also choose which potential elements of context to attend to and interpret with and which ones to ignore. In an ideal, cooperative conversation, the speaker and hearer will *harmonize* their contexts [*cf* Regoczei and Hirst 1990], *negotiating* what they deem to be relevant; but they're under no special obligation to do so.

In his fascinating book *The flight from ambiguity* [1985], Donald N. Levine points out that an aversion to ambiguity in communication, and hence to the kind of discretionary interpretation that I've just described, is a modern Western phenomenon, "unique" in world history [p. 21], whereas "most if not all of the literate civilizations have considered the cultivation of ambiguous locution to be a wonderful art" [p. 21–22]. For instance, in Somali,

> political arguments and diplomatic messages take the form of alliterative poems, mastery of which is a key to prestige or power. These poems typically begin with long, vague, circumlocutory preludes, introducing the theme at hand, which is then couched in allegory. … 'A poetic message can be deliberately misinterpreted by the receiver, without his appearing to be stupid. [The receiver may] go into further allegory, circling round the issue in other ways, to prevent direct confrontation.' [p. 23–24, quoting Laitin 1977, p. 39].

This approach to communication reaches its zenith in Amharic, of which

> one is considered a master … only when one's speech is leavened with ambiguous nuances as a matter of course. [p. 25]

Amharic conversation tends to be evasive; a speaker might say "Give me!" without any specification of what he wants; and even if asked, might not say what he really intends; nor would he be believed if he did [*ibid.*].

> When [an Amhara] talks, his words carry *double entendre* as a matter of course; when he listens, he is ever on the lookout for latent meanings and hidden motives. [pp. 27–28]

In other words, both the speaker and the interpreter have some discretion in *choosing* and *constructing* the context in which the interpretation is to be built. And while this discretion might be greater, and more explicitly licensed, in Amharic or Somali than in English, it's true in English too, even if our politicians don't routinely speak in alliterative allegorical poetry. We see it every day in our ordinary conversations [Devlin and Rosenberg 1996, p. 18], in advertising, political discourse, poetry, humor, allusion, persuasion and deception, negotiation of meaning, and misunderstanding and its repair [Hirst, McRoy, Heeman, Edmonds, and Horton 1994].

Here's a very simple example: In his first U.S. presidential campaign, in 1992, H. Ross Perot said, in reference to possible free trade with Mexico,

> [If] you don't care about anything but making money, there will be a giant sucking sound going south.[1]

Perot's remark was widely reported, and for many years after, even in 1999, was frequently alluded to by other speakers. Here is an example from 1995:

> The 'giant sucking sound' is not merely American jobs going to NAFTA and GATT cheap labor zones. The giant sucking sound is that as jobs and education diminish, our youth are being sucked into the jail industrial complex.[2]

Neither Perot nor NAFTA had been previously mentioned in the article, but the writer (Rev. Jesse Jackson) could allude to them anyway by this phrase, and *make* them part of the context. A quick Web search easily finds hundreds of such allusions to this one phrase. To fully understand the speaker's intent, the interpreter has to recognize the allusion and adjust the context of interpretation accordingly.

But of course, having discretion in constructing context does not mean having complete freedom. Obviously, considerable constraints arise from the message that is to be communicated, the circumstances under which the communication occurs, and the mechanisms of language itself. Even in Somali, the interpreter has to take the utterance itself as a given. And any language imposes rules as to how anaphora, for example, are to be interpreted with respect to the preceding text, and so any preceding text is necessarily an element of the context. Some aspects of the situation are also obligatorily included. For example, Deborah Tannen [1990] has shown that in American English, there are classes of sentences for which the gender of the speaker determines pragmatic aspects of the intended univocal interpretation, and the speaker's gender is therefore an element of the context that the interpreter ignores at his or her peril.

So what can actually be used as an element of context in natural language? Many other people have already offered inventories or taxonomies of the kinds of things that a speaker or interpreter must or may include in a context, and so I don't want to spend a lot of space here going into details. I need only point out that it includes just about anything in the circumstances of the utterance, and just about anything in the participants' knowledge or prior or current experience [*cf* Empson 1953]. So, Sperber and Wilson [1986] have argued in detail that a speaker or listener can use any fact or belief about the world that they have as an element of context. Ferrari

[1] U.S. Presidential Debate, Richmond, Virginia, 15 October 1992. *http://www.usia.gov/usa/infousa/facts/democrac/73.htm*

[2] Jesse Jackson, "32 years later: The dream unfulfilled." *JaxFax*, **3**(34), 8 August 1995. *http://www.cais.com/pcedge/test/rb/fx50824.html*

[1997] emphasizes the multimodal aspects of communication; he divides elements of context into those of linguistic context, perceptual context, intentional context, and encyclopedic knowledge, and he includes the message itself, along with all the circumstances of its utterance, in what he calls the "communicative situation". Zarri [1995] distinguishes between the *a priori* "internal / static" context and the "external / dynamic" context. The former includes knowledge of the language itself, such as the lexicon. Akman [1997] reiterates seven dimensions of context from the work of Wendell Harris, and emphasizes a view of context as a social construct. Manfred Pinkal [1985, p. 36] summarizes it all rather well:

> Aside from the surrounding deictic coordinates, aside from the immediate linguistic cotext and accompanying gestural expressions at closer view, the following determinants can influence the attribution of sense: the entire frame of interaction, the individual biographies of the participants, the physical environment, the social embedding, the cultural and historical background, and—in addition to all of these—facts and dates no matter how far removed in time and space. Roughly speaking, 'context' can be [*I'd rather say "draw on"—G.H.*] the whole world in relation to an utterance act."

So the discretion exercised by a speaker or an interpreter in constructing a context is, in effect, a determination of what, among all this, is and isn't *relevant* to the utterance that is to be interpreted [Sperber and Wilson 1986]. But this leads us to a terminological difficulty. For something to even be *considered* for possible relevance seems to imply that, regardless of the actual decision, that thing *is* an element of the context—otherwise, how could it come to be considered? There are *two* intuitive notions of 'context' here. The first, which I've tacitly been using up to this point, is the set of things that are used to build the interpretation with [Sperber and Wilson 1986, p. 15]; and the other is all that, *plus* the things that *aren't* used but nonetheless had a potential to have been used. For example, you might say that because the gender of the speaker is *sometimes* a factor in interpretation, it is therefore *always* an element of context, even if the interpreter doesn't always choose to use it. But then, by the same argument, you'd have to say that Ross Perot's giant sucking sound is always in the context because a speaker can always make an utterance that alludes to it. And by a similar argument, *everything* is in *all* contexts—much as Pinkal suggests. But that's not a very helpful view. A middle ground, and one that I lean toward, is to say that context involves a notion of *attention* to account for things that are at least considered for use in constructing the utterance or interpretation; to decide that something is *not* to be used in forming an interpretation is, in a sense, to use it in forming that interpretation!

We see then that context is both a psychological construct and, as Akman says, a social construct, and it's a social construct both in the sense that it is a construct *of* society and in the sense that it is constructed socially—in all our communication and social interactions—and constructed dynamically. It's not just a matter of moving axioms between pre-defined contexts with pre-defined lifting axioms.

Given all this, research on context in natural language starts to look quite familiar. In fact, much (or maybe most) research in natural language in artificial intelligence for the last 25 years and more can be seen simply as attempts to characterize context [*cf* Sowa 1995]. Roger Schank's scripts [Schank and Abelson 1977, Schank and Riesbeck 1981] and Gary Hendrix's partitioned semantic nets [1975] in the 1970s; my own marker passing in knowledge bases in the 1980s [Hirst 1987]; present-day statistical approaches based on lexical co-occurrence; my own group's recent use of *lexical chains* as "cheap" context for tasks such as segmenting discourse, finding real-word spelling errors, and automatically creating hypertext [Morris and Hirst 1991, Hirst and St-Onge 1998, Green 1997]—these are all really just attempts to provide or construct contexts with which utterances can be interpreted.

These approaches have had varying degrees of success. Some were simply wrong—that is, they made observably false assumptions about the nature of language. For example, Schank's scripts assumed that situations always uniquely pre-determine the word meanings and inferences that are applicable in the situation. McCarthy and Buvač's approach seems to be, in effect, Schankian. Buvač [1996], for example, chooses between two homonymous meanings of the word *bank* in a logical form based on the sentence *Vanja is getting money at a bank* by assuming that all other words in the sentence are unambiguous and can be used to find the exact right axiom in the commonsense context. As far as I can tell, the example relies crucially on the assumption of univocality of the words *get*, *money* and *at*, and if *at* were changed to *from*, the method would fail on the resulting polysemy or metonymy.

Other AI approaches to interpretation in context were perhaps a little more correct in principle, but still made either unrealistic assumptions about language or impractical assumptions about the knowledge sources upon which they were supposed to draw; my own work [Hirst 1987] on using semantic associations in a knowledge base as a context for disambiguation should probably go in the latter category. Nonetheless, this approach at least had the merit that interpretation was incremental, including the construction of context. It did not assume that parsing, let alone the building of a logical form, can occur prior to any consideration of context or to processes of disambiguation and interpretation.

So there's a sense in which just about all research in AI on natural language *is* research on context. And as we now see, it's rather different from context in KR.

There's one obvious objection to making this distinction between context in NL and context in knowledge representation and reasoning. Proper interpretation of natural language, we've been told for many years, *requires* knowledge representation and reasoning. So we'd better have a single theory of context that covers them both.

I have two responses to this. First, despite Sperber and Wilson [1986], it's becoming clear that while the knowledge used in interpreting natural language is broad, the reasoning is shallow. Although we can't yet characterize it precisely, it seems to be pretty much limited to reasoning about quite simple commonsense knowledge, knowledge of kinds, of associations, of typical situations, and even typical utterances. We don't do arbitrary reasoning in interpretation. So we don't need a very general theory.

Second, we probably *do* do arbitrary reasoning on arbitrary knowledge when we *assimilate* interpretations—when we build and refine our mental models. But there's no reason to think that that necessarily requires the same representations or mechanisms as are used in creating the interpretation. On the contrary, it's now clear that the mind uses many *different* kinds of representations and mechanisms. And, of course, to the extent that using natural language and representing and reasoning about knowledge are both cognitive activities, there's no reason to think that they are characterized by *any* AI-style formalizations—and plenty of reasons to think that they aren't, as Lakoff and Núñez [1997, pp. 22–23] have argued.

5 CONTEXT AS A SPURIOUS CONCEPT

So far, I've argued that the notion of 'context' can be defined only in terms of its effects in a particular situation. Just as a stressor is anything that stresses, one way or another, in at least one situation, context is something that constrains, one way or another, in at least one situation. In the case of natural language, many different kinds of things can be elements of context. Context in natural language is *constructed*, in part, by the speaker and the interpreter—it's not the same as context in knowledge representation.

In this light, context *simpliciter* can be seen to come dangerously close to being a spurious or incoherent concept in much the same way that 'absolute motion' is [Peacocke 1992]. In fact, there are a number of parallels between the two. In both cases, we have an intuition about the putative concept, and a very robust intuition at that. In our daily lives, we use what seems to be the notion of absolute motion in our navigation and moving around, and, as high-school science teachers know, it's not an easy notion to break away from. For obvious cognitive reasons, it's a psychologically compelling idea. We only need to look at the history of science to see how reluctantly it was given up, even by highly educated people. Yet now, a hundred years after the Michaelson–Morley experiment, it seems so obvious that ‘absolute motion’ is an incoherent or spurious concept that it’s hard to imagine how people ever thought otherwise.

Likewise, we use what seems to be a general notion of context when we build our interpretations of everything in our daily lives; but this is nothing more than an illusion that arises from our inability to examine our own mental processes of reasoning and language interpretation. Matthew Dryer [1997] has shown that the idea of sentence topic in linguistics, which has long been thought to be an intuitively well-founded concept, is in fact a chimera—what Dryer calls a “metalinguistic illusion”. Dryer has shown that just because a sentence is about something, it doesn’t follow that there’s any constituent in the sentence that’s what the sentence is about. All there is is discourse topic, even if the discourse is just a single sentence. Context *simpliciter* might turn out to be like this—seemingly intuitively well-founded, but revealed as a chimera upon deeper analysis.

And both absolute motion and context *simpliciter* are easy to formalize. Cartesian coordinates work quite nicely for the former in simple everyday applications. For the latter, McCarthy and Buvač’s [1997] formalization of context *simpliciter* can, under certain assumptions, find the price of airplane parts and disambiguate two homonymous senses of the word *bank* [Buvač 1996]. But however useful they are in local human day-to-day navigation, Cartesian coordinates are not a very useful formalization for what is now known in theoretical physics about the nature of space and time. And simple formalizations of context *simpliciter* might work on toy examples, but there’s no reason to expect them to apply to real-world natural language. On the contrary, a little analysis of what ‘context’ actually is suggests that they won’t.

6 CONCLUSION

Artificial intelligence in general is sometimes just a bit too impetuous in its desire to formalize things, and it tries to turn things into systems or logics without fully understanding them, as if simply by doing so they would thereby come to be understood. Sometimes this works; and sometimes it just leads to meaningless, ungrounded formal systems [Lakoff and Núñez 1997]. To someone with a hammer, every screw looks like a nail. And topics that deal with language, cognition, and acting in and interpreting the world get more than their share of this bad treatment.

This seems to arise from a combination of overenthusiasm for Western scientific method and a misunderstanding of the nature of language that sometimes seems to border on fear. In this view, language is a messy and highly imperfect medium that is not to be trusted, but rather must either be sidestepped entirely or be beaten into submission by means of logic and formalism. This is pretty explicit in the work of Frege and Bertrand Russell [1918, p. 205], for example. Maybe that's why Russell looked up to Wittgenstein, who engaged the difficult questions of language that Russell avoided and found some frightening answers—that some concepts can't be defined by necessary and sufficient conditions, for example. But all AI knows how to do is carry on as if Wittgenstein had never existed; nor Heidegger and Gadamer; nor Donald Levine; nor Sperber and Wilson; nor George Lakoff; nor Herbert Clark; nor Roy Harris; nor Harvey Sacks and Emanuel Schegloff and Harold Garfinkel and Erving Goffman. And I carry on that way too, at times—but at least *I* feel guilty about it.

In this paper, I've been rather negative and pessimistic in places, and I don't want to close on that kind of a note. After all, one thing that the field of artificial intelligence has certainly succeeded in over the years is expressions of unbounded optimism. So I want to close by emphasizing that we *do* have a good chance of getting a handle on 'context'—but we need to avoid premature, uninformed formalization. Situation theory [Devlin 1991] seems to me to be one especially good candidate. There is a strong intuitive relationship between the ideas of 'context' and 'situation'; situation theory has been under development for many years; and computational and linguistic concerns have been there from the start [Barwise and Perry 1983]. It is heartening to see books such as that of Devlin and Rosenberg [1996], who apply situation theory to real language in use and who say in their preface that their greatest intellectual debt is to Harvey Sacks. Work on formalizing context that uses situation theory, such as that by Akman and Surav [1996, 1997] and Ferrari [1997], is likely to be pointing us in the right general direction. There are also many other approaches to context—as can been in recent workshop proceedings on the topic [Iwańska 1995, 1997; Buvač 1995]—that offer considerable promise.

REFERENCES

Akman, Varol, 1997. "Context as a social construct." In Iwańska 1997, 1–6.

Akman, Varol and Surav, Mehmet, 1996. "Steps toward formalizing context." *AI Magazine*, **17**(3), Fall 1996, 55–72.

Akman, Varol and Surav, Mehmet, 1997. "The use of situation theory in context modeling." *Computational Intelligence*, **13**(3), August 1997, 427–438.

Barwise, Jon and Perry, John, 1983. *Situations and attitudes.* Cambridge, MA: The MIT Press.

Buvač, Saša (editor), 1995. *Working notes, AAAI Spring Symposium on Formalizing Context*, Stanford University.

Buvač, Saša, 1996. "Resolving lexical ambiguity using a formal theory of context." In: van Deemter, Kees and Peters, Stanley, 1996. *Semantic ambiguity and underspecification*, Stanford, CA: CSLI Publications. 101–124.

Devlin, Keith, 1991. *Logic and information.* Cambridge University Press.

Devlin, Keith and Rosenberg, Duska, 1996. *Language at work: Analyzing communication breakdown in the workplace to inform systems design*. Stanford: CSLI Publications.

Dryer, Matthew, 1997. "The myth of sentence topic." Unpublished.

Empson, William, 1953. *Seven types of ambiguity*, third edition. London: Chatto and Windus.

Ferrari, Giacomo, 1997. "Types of context and their role in multimodal communication." *Computational Intelligence*, **13**(3), August 1997, 414–426.

Green, Stephen, 1997. "Building hypertext links in newspaper articles using semantic similarity." *Third Workshop on Applications of Natural Language to Information Systems (NLDB '97)*, Vancouver, June 1997, 178–190.

Hendrix, Gary G., 1975. "Expanding the utility of semantic nets through partitioning." *Advance papers of the 4th International Joint Conference on Artificial Intelligence*, 115–121.

Hirst, Graeme, 1987. *Semantic interpretation and the resolution of ambiguity.* Cambridge University Press.

Hirst, Graeme, 1995. "Near-synonymy and the structure of lexical knowledge." *Working notes, AAAI Spring Symposium on Representation and Acquisition of Lexical Knowledge: Polysemy, Ambiguity, and Generativity*, Stanford University, March 1995, 51–56.

Hirst, Graeme and St-Onge, David, 1998. "Lexical chains as representations of context for the detection and correction of malapropisms". In: Fellbaum, Christiane (editor), *WordNet: An electronic lexical database and some of its applications*, Cambridge, MA: The MIT Press, 1998. 305–332.

Hirst, Graeme; McRoy, Susan; Heeman, Peter; Edmonds, Philip; and Horton, Diane, 1994. "Repairing conversational misunderstandings and non-understandings." *Speech communication*, **15**(3–4), December 1994, 213–229.

Iwańska, Łucja (editor), 1995. *Working Notes, Workshop on Context in Natural Language Processing, 14th International Joint Conference on Artificial*

Intelligence, Montreal.

Iwańska, Łucja (editor), 1997. *Working notes, AAAI Fall Symposium on Context in Knowledge Representation and Natural Language*, Cambridge, MA.

Laitin, David D., 1977. *Politics, language, and thought.* The University of Chicago Press.

Lakoff, George and Núñez, Rafael E., 1997. "The metaphorical structure of mathematics: Sketching out cognitive foundations for a mind-based mathematics." In: English, Lyn D. (editor), *Mathematical reasoning: Analogies, metaphors, and images*. Mahwah, NJ: Lawrence Erlbaum Associates. 21–89.

Levine, Donald N., 1985. *The flight from ambiguity: Essays in social and cultural theory*. The University of Chicago Press.

McCarthy, John, 1987. "Generality in artificial intelligence." *Communications of the ACM*, **30**(12), 1030–1035.

McCarthy, John, 1996. "A logical AI approach to context." Unpublished note, 6 February 1996. *http://www-formal.stanford.edu/jmc/logical.html*

McCarthy, John and Buvač, Saša, 1997. "Formalizing context (expanded notes)." In: Aliseda, Atocha; van Glabbeek Rob; and Westerståhl, Dag (editors), *Computing Natural Language*. Center for the Study of Language and Information, Stanford University.

Morris, Jane and Hirst, Graeme, 1991. "Lexical cohesion, the thesaurus, and the structure of text." *Computational linguistics*, **17**(1), March 1991, 21–48.

Peacocke, Christopher, 1992. *A study of concepts*. The MIT Press.

Pinkal, Manfred, 1985. "Kontextabhängigkeit, Vagheit, Mehrdeutigkeit." In: Schwarze, Christoph and Wunderlich, Dieter (editors), *Handbuch der Lexicologie*, Königstein: Athenäum Verlag. The quotation used is translated and cited in: Quasthoff, Uta M. "Context." In: Asher, R.E. (editor) *Encyclopedia of Language and Linguistics*, Pergamon Press, 1994, 730–737.

Regoczei, Stephen and Hirst, Graeme. "The meaning triangle as a tool for the acquisition of abstract, conceptual knowledge." *International journal of man–machine studies*, **33**(5), November 1990, 505–520.

Russell, Bertrand Arthur William, 1918. "The philosophy of logical atomism." In *The philosophy of logical atomism and other essays 1914–19*, edited by John G. Slater (The collected papers of Bertrand Russell, volume 8), London: George Allen & Unwin, 1986. 157–244.

Schank, Roger C. and Abelson, Robert P., 1977 *Scripts, plans, goals, and understanding.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Schank, Roger C. and Riesbeck, Christopher K., 1981 "The theory behind the programs: A theory of context." In: Schank, Roger C. and Riesbeck, Christopher K. (editors), *Inside computer understanding: Five programs plus miniatures*, Hillsdale, NJ: Lawrence Erlbaum Associates.

Shoham, Yoav, 1991. "Varieties of context." In: Lifschitz, Vladimir (editor), *Artificial intelligence and the mathematical theory of computation: Papers in honor of John McCarthy*, Academic Press. 393–407.

Sowa, John F., 1995. "Syntax, semantics, and pragmatics of contexts." In Iwańska 1995, 145–154.

Sperber, Dan and Wilson, Deirdre, 1986. *Relevance: Communication and cognition*. Harvard University Press.

Tannen, Deborah, 1990. *You just don't understand: Women and men in conversation.* New York: William Morrow and Company.

Zarri, Gian Piero, 1995. " 'Internal' and 'external' knowledge context, and their use for the interpretation of natural language." In Iwańska 1995, 180–188.

Graeme Hirst's recent research topics have included near-synonymy and measures of semantic distance. He presently serves as book review editor of *Computational Linguistics*. Hirst's address is: Department of Computer Science, University of Toronto, Toronto, Ontario, Canada M5S 3G4; e-mail: *gh@cs.toronto.edu*; URL: *http://www.cs.toronto.edu/~gh*

An earlier version of this work was presented in an unpublished invited talk at the AAAI Fall Symposium on Context in Knowledge Representation and Natural Language, Cambridge, Massachusetts, 8 November 1997. I am indebted to Łucja Iwańska for inviting me to present that talk; and to Stephen Regoczei for enabling me to write it, by means of many discussions over the years on many of the issues herein. I was also helped by discussions with, or provocations from, Nadia Talent and Chrysanne DiMarco, and by reading the work of Varol Akman and Bruce Edmonds. I am grateful to my fellow Canadian taxpayers for a research grant from the Natural Sciences and Engineering Research Council.