

Indexicality and Representation

Peter Bosch
Institute for Knowledge Based Systems
IBM Germany Scientific Center
Schloßstr. 70, D-7000 Stuttgart 1
bosch@dsØlilog.bitnet

Abstract

In the vast majority of cases linguistic comprehension does not care much about what there is in the world and what state the world is in. Comprehension is often satisfied with satisfaction: you have understood somebody's statement once you have set up a model that satisfies the statement.

But occasionally reality matters. When Albert confesses to Bob: "I'm having an affair with your wife", and Bob's comprehension goes no further than constructing a model that satisfies Albert's statement, i.e. a model in which a speaker is having an affair with his listener's wife, then there is nothing that could make Albert nervous or that could make Bob jealous. Merely constructing a model does not link the statement to the world.

But how do we model this step from a model to real life, from "I" or "you" to a real person? A proposal sketched near the end of this paper is that – despite first appearance – we don't and we needn't.

1. Introduction

It would seem obvious that there is an important difference between sentences like (1) and sentences like (2). The truth of (1) is independent of when and by whom it is uttered, whereas the truth of (2) depends crucially on who utters it and when. (The examples are Bar-Hillel's (1954:359)).

- (1) Ice floats on water.
- (2) I am hungry.

When our interest is in relations of logical consequence, cases like (2) are an unpleasant complication, certainly if the definition of logical consequence we want to give is a formal definition, and hence cannot rely on anything but the formal properties of those objects between which relations of logical consequence hold. English sentences that contain indexical expressions, like sentence (2), thus pose a *prima facie* difficulty. Indexical expressions, like "I", "you", "here", "now", etc. change their reference according to the situation in which they are used. Accordingly, also the truth and falsity of sentences that contain indexicals vary, depending on properties of the situation in which they are uttered, and hence such sentences are not suitable terms for relations of logical consequence.

But what are suitable terms? When what we are after is a semantics for English sentences then what we need is formal objects that stand in a regular relation to English sentences (at least one such object for each sentence) and that are invariant with respect to relations of logical consequence, as sentences like (2) plainly are not.

The problem at hand is not ambiguity, even though also ambiguities may be, and frequently are, resolved by consideration of the relevant context. By an ambiguous sentence I understand a sentence that is associated with a small and manageable number of different sets of truth conditions. Since ambiguity is due either to the occurrence of an ambiguous word or to a structural equivocation, it can be dealt with by fairly simple disambiguation techniques.

Ambiguous words, or so I shall assume for the purposes of this argument, are listed completely in the lexicon, and their different senses are formally distinguished by different subscripts added to their spelling. And sentences may be regarded as pairs of a sequence of characters (including subscripts, blanks, and punctuation marks) and a structural description. Indeed, the disambiguated sentences we now have are no longer strictly speaking English sentences. But they are formal objects, they are formally distinct from each other, and there is at least one of them for each English sentence.

But such disambiguation techniques, however adequate for ordinary ambiguity, would seem insufficient for context dependence proper. First, there is good reason to believe (cf. e.g. Ziff 1973, Travis 1977, Clark and Clark 1979, Bosch 1983 Ch.3, 1985a) that the number of contextually different word senses is potentially infinite or, at any rate, forbiddingly large, and prevents any form of listing.

Second, and independently: even if the number of word senses were manageable, the variation in reference we find with indexical expressions certainly is not. – So even if the differences and connections between run-of-the-mill ambiguity and context dependence are not clear in all respects, we have here at least a couple of technical reasons not to attempt a treatment of indexicality in terms of established disambiguation techniques.

But we may still try to keep our account of indexicality parallel to disambiguation if, in order to find objects that are invariant under changes of utterance context, we move from English sentences (or, better, disambiguated English sentences) to pairs of a sentence and a context of utterance. These new objects clearly satisfy the requirement of being invariant terms for relations of logical consequence in a way that disambiguated sentences are not. But there still remains one point of worry: can these new objects be construed as formal objects? Can each two such objects be distinguished on purely formal grounds?

For convenience of expression I shall refer to these pairs of a disambiguated sentence and a context as “C-propositions” – “propositions”, because they seem to be similar in some respects to what others mean by this term, and the “C” as a warning that the term is only used for convenience and no more should be read into it than what follows from our explicit statements. – One may note that C-propositions are particularly close to Stalnaker’s ‘propositions’ and to Kaplan’s sentence ‘contents’ (cf. Stalnaker 1972, Kaplan 1978, Bosch 1982). But I am not interested here in the peculiarities of these or other specific proposals and therefore prefer the non-committal term “C-proposition”.

Again the question: can C-propositions be construed as formal objects? Is there a way of distinguishing formally between each two C-propositions? – We saw above that sentences cause no problem as formal objects. A sentence may be construed as a sequence of characters; and each character may be identified with a set of inscriptions (cf. Quine 1960:195). – But how can contexts of utterance be construed in a similarly unobjectionable fashion? To this question I shall attend in a moment. But first I must broaden the issue a little.

2. Psychological Considerations.

Our interest in the notion of logical consequence with respect to the sentences of a natural language is due to the assumption that this notion plays a central role in an explication of the semantics of a natural language: what follows from the truth of a sentence *S* tells us a great

deal about what S means, at least what S means with respect to, or in terms of, other sentences. Since there is too much fluctuation with respect to logical consequence relations, and hence with respect to the meaning (in this specific sense) of English sentences, we must look for entities that are stable in this respect, for instance C-propositions. C-propositions may thus be regarded as meanings of disambiguated English sentences relative to a context, in the sense that they are equivalent to that part of sentence meaning that is invariant in relations of logical consequence, i.e. that portion of meaning that matters for the preservation of truth.

Understanding an utterance of an English sentence would then, according to the above considerations, be the same as grasping what C-proposition the sentence amounts to when uttered on the particular occasion in question. If this connection were non-existent, the investigation of relations of logical consequence would be irrelevant to an explication of linguistic comprehension.

Now there is the following difficulty. C-propositions are pairs of a sentence and a context. And we require that both sentence and context are formal objects. But do human speakers, when they understand a sentence in a context, avail themselves of anything that would amount to a formal representation of the context? – In practice this requirement would mean that speakers should be able to ‘transcend’ their utterance contexts, as it were, to make explicit for themselves what aspects or features of the context are influencing their understanding of the sentence, or to paraphrase each indexical sentence by an equivalent non-indexical sentence. – But ordinary comprehension processes, or so it seems, do not care about such a requirement.

Consider a situation where both speaker and listener neither know the time, nor the place at which they are – two Rip van Winkles, if you like, who not only have slept for an indefinite period of time but who have also been sleep-walking and find themselves in an unknown place. They will, in such a situation, still use indexical expressions that depend on their own current spatio-temporal position. And not only is the functioning of these expressions entirely unproblematic here, but it is – given that neither of them know the place or time – the only way available to them for talking about their spatio-temporal position. – They could not translate their indexical sentences into non-indexical ones, have no way of ‘transcending’ their situation, and have no representation of their utterance context except their own current and indeed very limited one.

Since human speakers ordinarily communicate by means of indexical sentences without any difficulty of a principled nature, the results of their comprehension processes, at least in these cases, cannot be C-propositions, i.e. representations that make context fully explicit and that are, in this sense, context independent, or non-indexical, representations. The mental representations that result from human comprehension plainly do not contain any information the listener does not have, and if he has no context-independent way of describing, say, the place or the time of the communication situation, some equally indexical equivalent of “here” and “now” will inevitably form part of his mental representation of what he has understood.

In brief, then, the argument runs as follows: relations of logical consequence require terms that are invariant with respect to context; such terms (C-propositions) cannot play a decisive role in comprehension processes, because natural language communication does not require the availability of context-independent representations. Hence: as far as the human processes of linguistic comprehension are concerned, relations of logical consequence cannot play a decisive role either.

3. Two Problems

I have sketched two problems. The first of a general philosophical nature and concerned with the very possibility of a formal reconstruction of context, and the second of a psychological kind and concerned with the role of explicit context representation in human comprehension. While the second problem is taken very seriously, by at least some semanticists (cf. Partee 1980 and references there), the first, philosophical, problem has been given next to no attention. The general position seems still very close to that of Bar-Hillel (1954:369) in the very early days of the semantics of indexicality:

“Not in every actual communicative situation could every indexical sentence be replaced, without loss of information, by a non-indexical sentence; but there is, on the other hand, no indexical sentence which could not be replaced by a non-indexical sentence, without loss of information, in some suitable communicative situation”.

While Bar-Hillel acknowledges our second problem, he takes a traditional logical empiricist stand with respect to the first. He explicitly endorses Carnap’s view that “non-indexical languages are sufficient for the formulation of any body of knowledge” (1954:367).

For the remainder of this paper I shall be concerned with this notion of non-indexical representation. I shall say no more on the psychological question. That, in fact, in human comprehension processes C-propositions and their like do not play a significant role is something I take for granted. But this does not imply that C-propositions are thereby uninteresting for the semanticist. They may well play their part in a philosophical model for human understanding. And a model as such – philosophical or in the empirical sciences - may be adequate, in the sense that it helps us understand certain matters better, even though some parts of the model may stand in no clear relation of correspondence to anything in the object domain which the model is a model of.

4. Non-indexical representation

In order to make precise sense of the notion of a C-proposition along the lines suggested above, and for the purposes suggested above, we need not only a formal notion of a (disambiguated) sentence, but we require an equally formal representation of contexts of utterance.

This requirement may also be formulated by saying that we need something like a type-token distinction for contexts and, with it, a formal definition of the notion of a context type. We are surely not interested, neither as speakers of a language nor as semanticists, in each and every possible difference that there may be between any two actual speech situations. Most differences between real contexts of utterance do not matter for the C-propositions resulting from those contexts, i.e. they could not influence logical consequence. – And this is just as well, because contexts, like any other concrete entities, cannot have complete and finite descriptions. If each and every difference were to matter, the entire enterprise would be doomed to fail from the very start. – What remains troublesome though is the fact that it is not always the same features of real utterance situations that matter for the C-propositions expressed in those situations. If, for instance, I say “I love you”, it is usually taken to be sort of important who I am speaking to, but place and time, or climatic conditions and the rheumatism of my neighbour, for that matter, don’t seem to contribute much. But when I say “I’ve never been here before”, time and place of utterance become important, while we can’t care less about who I am speaking to. And, eventually, also climatic conditions and my neighbour’s rheumatism may occasionally play their part, for instance for the understanding of an utterance of “He’s always bad in this weather”.

This fluctuation in what is and what is not relevant for whether two contexts count as the same would not be so much of a problem if, at least, we had some assurance that the number of parameters that are relevant for the identification of context types would be limited. This, however, is not the case. Lewis (1981:87) acknowledges the difficulty, and so does Cresswell (1973:111f). Cresswell seems convinced, though, that the sentence at issue may help select those parameters that are relevant for the individuation of contexts with respect to that sentence. As long as there are explicit indications in the sentence as to the relevant parameters, such as tensed verbs, adverbs of place and time, or demonstrative pronouns, Cresswell is certainly right. But how is the sentence going to tell us, for instance, that the reading of one of its predicates varies crucially with the context, like e.g. “hexagonal” in the sentence “France is hexagonal” (cf. Lewis 1981)? Two contexts of utterance that differ along the lines along which typical geographical and typical geometrical contexts differ would in this case have to count as different contexts. If one of the preceding sentences is “Italy is sort of boot-shaped”, the corresponding reading of “hexagonal” should make our sentence come out true. But if one of the preceding sentences is “Geometrical shapes are ever so much simpler than shapes in geography”, our sentence is less likely to count as true (cf. also Clark and Clark (1979) for a discussion of cases that are probably far more troublesome). – The problem is not that the shape of France has changed, but that the two contexts of utterance must count as different for the purposes of interpreting our sentence.

Or what if the interpretation of a sentence varies with the context of utterance, without there being any particular expression in the sentence that would be responsible? Take a sentence like “There is coffee in the hall”, either as an announcement of a coffee break during a conference, or as an English translation of the triumphant statement of a 19th Century Prussian tax officer (“Kaffeeschnüffler” – as they were called) who has just discovered evidence, in the form of two coffee beans under a rug, of what was probably illegal and untaxed imports. The sentence can hardly be defended as true in the first utterance situation if the fact of the matter consists of those two coffee beans under the rug, whereas no-one would convince our tax man that two beans should be insufficient to make his statement true. – Examples of this kind, where the interpretation of the sentence is clearly context dependent but does not seem to hinge upon any particular expressions in the sentence are easily multiplied (cf. Ziff 1973, Bosch 1985a). And even if we force the responsibility for the context dependence of such sentences on particular words, in the case at hand upon “there is” – perhaps already compositionally limited by its combination with “coffee” – I still fail to see how this could significantly limit the number and kind of parameters that would enter the truth conditions of our sentence.

Hopes for a formal definition of the notion of context type are not flying high when we think of such cases. – And why maintain hope artificially where there is none? I think we can be more radical. Let me propose the following argument. Cresswell is clearly right when he says that the identity of contexts of utterance is influenced by the sentences whose contexts they are. But this accounts only for one set of factors. Whether or not two real contexts of utterance may count as identical, i.e. as of the same context type, also depends on preceding sentences (most obviously, but not only, in cases where the sentence under consideration contains anaphoric expressions that are coreferential with expressions in preceding sentences) and – via them (or directly) – on *t h e i r* contexts, plus, of course, on any amount of situational context not mediated via sentences. How then are we to give a formal definition of context types if the semantically relevant context type depends itself on the consideration of contexts? All we get is recursion.

If this argument is correct, we may wonder about the roots of the identity problem for context types. There seemed to be no such problem with regard to sentence types, or types of

disambiguated sentences. And there are infinitely many sentence types just as there are infinitely many context types. How and why are contexts different from sentences?

The crucial difference is that the set of sentences of a natural language is finitely differentiated (cf. Goodman 1969:135f), whereas the set of utterance contexts is not.

For each two sentence types, S_i and S_j , and each sentence token s that does not actually count as a token of both types, it is at least theoretically possible to decide either that s is not an instance of S_i or that s is not an instance of S_j .

There is, however, no such assurance for contexts of utterance. There is no such thing as a fixed set of context types in the sense as there is a set of sentence types; in other words: there is no way of knowing what may count as a relevant difference between two context tokens. Or, if we already assume a parametrization of contexts, i.e. a (possibly infinite) list of all those features that may, at least occasionally, count as relevant distinctions between contexts, then there is still no way of knowing whether or not the presence or absence of a particular feature will matter for the identity of any particular context on a given occasion. One of the causes for this problem is that a particular feature of a particular context may count as relevant when the context is approached from one particular preceding context and may not count as relevant when the context is approached from a different context. – And this consideration obviously invites recursion.

At this point the theoretician may be tempted to look to empirical constraints for help: are there not limitations to human perception that could rule out certain differences as irrelevant, thus putting a stop to recursion? In a sense, yes. There are fairly clear physiological limitations. But no such limitation can have an absolute character, because perception is an intelligent process, far from purely passively receptive. Consider the following Gedankenexperiment, just as an illustration of this more general point by means of a comparatively simple case from visual perception.

Suppose we have a section from the colour spectrum, which we cut into a few thousand or more narrow strips so that we are unable to distinguish the colour of any two neighbouring strips, say A and B in Fig.1. In other words, if our context consists of A and B only, A and B are indistinguishable. Now we add to our context B's other neighbour, C. And suppose that B and C are as indistinguishable from each other as are A and B, while A and C are distinguishable. (I suppose it is reasonable to assume that we can find a way of cutting up the spectrum so that this is actually the case). Now, from what we have observed, indistinguishability of A and B and of B and C, plus distinguishability of A and C, the reasonable conclusion is that A, B, and C must all be different in colour. In other words, in the first context (which consists of A and B only) A and B are indistinguishable, and in the new context, with C added, A and B are distinguishable. The point is not just the familiar one that perception, even here on the very simple level of perception of colour differences, is partly a matter of inference. The more important aspect of our thought experiment is that even below the level of visible difference empirical identity relations may change when the set of contrasting objects, i.e. the context, is changed.




Fig.1

Let me generalize this case. Suppose we have a context, c_1 , in which two objects, a_1 and a_2 , count as the same thing in the sense of being two tokens of the same type. Either because they are, as in the case just considered, actually indistinguishable in c_1 , or because, in c_1 , we cannot, or have no reason to, attach any significance to any differences between them. Both a_1 and a_2 then are taken as tokens of one type, say the type A.

Now suppose the context c_1 is changed to a context c_2 by the addition of further objects, among them a_3 , which is not of type A, but is seen as a token of a further type, A'. It is quite conceivable now that, in c_2 , where both A and A' are relevant, we would regard a_2 as of A', as well as of A, whereas a_1 clearly does not form an instance of A'.

Consider for instance a context in which the two objects in Fig.2 count as two instances of the same type, differing only in orientation (which is discounted in the context at hand as irrelevant). This is a natural view in contexts where objects can be turned around, or mirrored, with no restrictions. Now consider Fig.3. The addition of the new, third object, may invite an overall view of the three objects as, say, printing characters. And for printing characters, obviously, orientation matters, and the identity between our first two object has gone. The second and third object however, may now perhaps both be regarded as instantiations of logical 'and'-signs, whereas the the first would count as an 'or'-sign. The first one is then the odd one out, while the context allows us to view the second and third as identical. Still, if we resist the interpretation of all three objects as printing characters and refuse to count orientation as relevant, then the original identity can be maintained and the third object is the odd one out.



Fig.2




Fig.3

Since perceptual or empirical identity in general can only rest on the irrelevance, or, in the extreme case, the absence, of perceptible difference, i.e. on indistinguishability, and since both relevance and distinguishability depend on the contrasts that are available in the context in which the comparison is made, this means that there is no context-independent way of identifying, and therefore, in the more general case, of classifying anything.

To the extent that contexts are objects like any others, these considerations of identity and classification also apply to contexts. In practice, this can mean either that the identity of a context from one occasion to the other, or the classification of a context as a particular kind of context, depends – in theory – on some wider context within which the question of identity of contexts arises. Or it may mean that identity depends on a preceding context from which we enter a particular new context that must be classified.

Now, plainly, the upshot of all this is that identity and classification of contexts, as it is obviously crucial for linguistic comprehension – i.e. for the identity of sentence-context pairs (C-propositions) – is possible only with respect to some context that is, at least for the time being, taken as primitive. And this is nothing but a generalization of Karl Bühler's (1978:107) notion of an "I-Here-Now"-Origo, the origin of a speakers momentary system of cognitive coordinates.

Let me stress that the latter point is not first and foremost a matter of any particularly human limitations of perception, but is a more general matter of the partiality of experience. It is indeed possible to define identity proper on the basis of indistinguishability:

For any x and any y : x is identical to y iff
there is no z are such that z is indistinguishable from x but not from y .

For any partial experience (and thus all human experience), however, this definition is obviously unavailable.

5. Formal semantics

Formal semantics of natural language is concerned with an explication of semantic notions with respect to natural language. On the most general level, the goal is a characterization of truth and logical consequence for natural languages. This task, according to our argument up to now, crucially involves a formal reconstruction of the notion of a context.

However, as we saw, this notion poses insurmountable problems for a formal reconstruction and thereby endangers the whole enterprise: no formal reconstruction of context, hence no C-propositions, hence no characterization of truth and logical consequence for natural languages – and hence no formal semantics of natural language.

But we are not quite that far yet. In the above exposition, I have followed a traditional – and as now appears, self-destructive – line of thinking in formal semantics, even at points, where there is no need to do so. Most importantly, I have uncritically accepted the implicit notion that relations of inference, as they are essential for natural language semantics, are relations of logical consequence. But is this really so? Would it not be sufficient to have inferences that operate reliably just within one particular context, plus, and this is of course the same, inferences that allow us the transition from one context to the other, provided there is a particular third context within which the first two contexts can be related to each other? In each particular case, we have then one context, the current context at hand, that is taken as primitive and that provides the required standards of identity for other contexts.

What we want then, instead, or rather on top, of a theory of logical consequence, is a theory of contextual inferencing. Contextual inferencing does not require non-indexical representations, because context is not allowed to vary. It does not matter who the referent of a particular token of “you” is (we don’t need a context-independent representation of its referent) as long as we have a mapping from “you” into the particular context within which the reference must be understood. It may happen that this is a mapping onto the same referent that, in this other context, would be referred to as “I”. – And what more does the listener need than the information that it is him who is being addressed (never mind a context-independent representation of his identity)?

Sure enough, such a theory of contextual inference is much less powerful than a theory of logical truth. Notions like logical necessity and possibility cannot be reconstructed in terms of contextual inference. Instead we get epistemological possibility and necessity. And this is probably all we need for natural language comprehension.

6. The philosophy of it

Near the beginning of this paper I assumed, following the tradition of logical grammar, that the explication of relations of logical consequence between sentences is a core area of semantics. The instability of actual sentences of natural language with respect to truth, however, forced us to look for other terms than sentences as the terms of consequence relations. Such terms would need to be formal objects, like sentences, but they have to be insensitive, in their truth conditions, to any changes of context. It turned out that no such context-independent, or non-indexical, representations can be constructed, due to the impossibility of a formal representation of context.

Does this mean that, once we decide to take indexicality and context serious, relations of logical consequence become irrelevant to the semantics of natural language? Certainly not. It only means that things get a little more complicated. Relations of logical consequence only hold inside one context and once we transgress context boundaries, we first have to reconstruct the preceding context in terms of the new, then current context before any logical inference can be applied. This reconstruction however is not itself a matter of pure logic, but makes use of non-logical rules and axioms, many of them with a default character. The reconstruction of one context from the point of view of another as a whole is not monotonous or straightforwardly incremental. In brief: contextual inferencing comes much closer to the processes that are investigated in AI under the label “reasoning” than they are to the comparatively simple uses of logic we are familiar with from the tradition of formal semantics of natural language. – The morale of my argument then is not that logic is irrelevant or less relevant to the semantics of natural language, but that the role of logic in the semantics of natural language is a more intricate one than we are inclined to tell our first year students, and than we happily admit to ourselves. – This of course is nothing new at all, but a mere triviality. Having understood the point, the real work is in working out the details. In this respect, the lesson from the above is that natural language semantics and reasoning (in the AI sense) have a good deal of their task in common.

References

- Bar-Hillel, Yehoshua (1954): “Indexical Expressions”. *Mind* 63:359-379
- Bosch, Peter (1982): “The Role of Propositions in Natural Language Semantics”. In: W. Leinfellner e.a. (eds.): *Language and Ontology. Proceedings of the 6th Wittgenstein Symposium. Hölder-Pichler-Tempsky. Vienna.*
- Bosch, Peter (1983): *Agreement and Anaphora. Academic Press. London.*
- Bosch, Peter (1985a): “Lexical Meaning Contextualized”. In: G. Hoppenbrouwers/ P. Seuren/ T. Weijters (eds.): *Meaning and the Lexicon. Foris. Dordrecht.*
- Bosch, Peter (1985b): “Propositionen”. In: Thomas Ballmer/ Roland Posner (eds.): *Nach-Chomskysche Linguistik. De Gruyter. Berlin.*
- Bühler, Karl (1978): *Sprachtheorie (first edn. 1934). Ullstein. Berlin.*
- Clark, Eve V. / Herbert H. Clark (1979): “When Nouns Surface as Verbs”. *Language* 55:767-811
- Cresswell, M.J. (1973): *Logics and Languages. Methuen. London.*
- Davidson, Donald / Gilbert Harman (eds.) (1972): *Semantics of Natural Language. Reidel. Dordrecht.*
- Goodman, Nelson (1969): *Languages of Art. Oxford University Press. London.*
- Kanger, Stig / Sven Öhman (eds.) (1980): *Philosophy and Grammar. Reidel. Dordrecht.*
- Kaplan, David (1978): “On the Logic of Demonstratives”. *Journal of Philosophical Logic* 8:81-98.
- Lewis, David (1981): “Index, Context, and Content”. In: Kanger / Öhmann (eds.)
- Partee, Barbara Hall (1980): “Montague Grammar, Mental Representations, and Reality”. In: Kanger / Öhmann (eds.)
- Quine, W.V. (1960): *Word and Object. MIT Press. Cambridge, Mass.*
- Stalnaker, Robert (1972): “Pragmatics”. In: Davidson / Harman (eds.)

- Travis, Charles (1977): "What Semantics Is Not". In: P.A.M. Seuren (ed.): Symposium on Semantic Theory. Dept. of Philosophy. University of Nijmegen.
- Ziff, Paul (1972): "What Is Said". In: Davidson / Harman (eds.)
- Ziff, Paul (1973): "Something About Conceptual Schemes". In: Glenn Pearce / Patrick Maynard (eds.): Conceptual Change. Reidel. Dordrecht.