

# CONTEXTS IN CONTEXT

Patrick J. Hayes  
IHMC, University of West Florida

## INTRODUCTION

This purpose of this note is lexicographic rather than theoretical: to tease apart some of the meanings of the word 'context' as it is used in the technical literature. I will distinguish four different kinds of context and the presuppositions that lie behind them. I hope to clarify what I believe to be a difference between two intellectual traditions, each of which brings a different collection of unspoken assumptions.

Speaking about unspoken assumptions is a tricky business. Readers who already possess one set of assumptions may find the business of explicating them tedious and elementary, while those who do not may find it ludicrous or incomprehensible. Some readers may interpret everything one writes in only one sense, thus rendering the discussion unintelligible. Also, a proper survey of these issues is beyond the scope of this short paper, so my accounts will be somewhat simplified and may seem like caricatures. If any of the following seems either pathetically obvious or outrageously incorrect, please consider the possibility that others may view things differently. I do not mean here to advocate any of these various opinions or methodologies in preference to the others, but merely to try to clarify these potentially misleading differences. The lexicographer's task is simply to report, not to judge.

## TWO APPROACHES TO LANGUAGE

Most of the literature on language agrees in broad terms with the following general picture. People use language to communicate, and have some kind of internal representation of what it is that they mean to communicate and understand others to be communicating to them. The external language is used for communication, while

the internal representation encodes beliefs and supports mental processing. Language comprehension and production establish connections between the external language and the internal representational code.

Unlike external language, the internal code cannot be publicly observed, so there are many ideas about what its structure might be. Almost the entire literature of psycholinguistics and NL work in AI is concerned one way or another with hypotheses about the internal code. Fodor argues that this 'language of thought' must have a productive syntax, but Johnson-Laird and others think of it as consisting of 'mental models' which are more similar to diagrammatic images. It might consist of many rather isolated systems (spatial, visual, etc.), or may all be representable in a single logical notation, as AI often assumes.<sup>1</sup> To avoid taking sides on any of these disputes, I will here use the deliberately artificial terms 'EL' to refer to the external language, i.e. the normal subject-matter of linguistics, and 'IC' to refer to the internal code, i.e. the internal mental representation of the result of the language comprehension process.

The first intellectual tradition, that of semantic linguistics, focuses on EL. Following the early work of Montague, its ambition is to provide a semantics for EL in the form of a model theory defined directly on the syntax of EL sentences provided by linguistic grammars. This semantics should fit naturally onto the syntactic structure assigned to sentences of EL by a grammatical

<sup>1</sup>More exactly, a logical notation is used to represent the content of the IC. From an AI perspective, this is the minimal hypothesis (at the 'epistemological level', McCarthy and Hayes 1969) and further details of representation types, modularization, etc., represent further hypotheses about how the mental content might be organized.

theory. The linguistic syntax is central to this tradition; for example, any systematic difficulty in making a semantics properly conform to syntactic structure is considered to be a criticism of the grammatical theory. On this view, the challenge provided by the contextual sensitivity of EL is to elegantly incorporate extra structure (representing the contextual important information) into the semantic theory so that the meaning of larger syntactic forms can still be regarded as somehow formed from those of their syntactic parts plus this extra structure, while conforming to the constraints of the currently accepted grammatical theory. Much of the modern work in linguistic semantics is concerned to identify suitable structural devices that preserve compositionality in this way, such as Kamp's 'discourse reference structures' (1990) and terMeulen's 'chronoscopes' (1995)<sup>2</sup>. Since all native users of language have reliable intuitions about correct and incorrect uses of NL sentences in various settings, these proposed theoretical accounts of contextual meaning can be tested against linguistic intuitions in much the way that judgments of syntactic naturalness are used to test proposed grammatical theories, fitting smoothly into the usual linguistic methodology. While Montague's original vision of a model theory for English now seems somewhat simplified, the central concern of this area is still a referential semantic theory which connects EL to the world it describes.

There are many differences within this large field of study, but all work in this tradition shares some common themes and attitudes. It has a central concern with the syntax of EL and uses grammatical criteria to isolate particular research topics, and is only peripherally concerned with language use. It has a methodological aversion to psychological theorizing or any concern with details of hypothesized internal codes

<sup>2</sup>It may be worth emphasizing, for readers in the AI tradition, that such structures are *not* thought of as data structures to be used in a computational process of comprehension.

or processes<sup>3</sup>. The inferences whose validity the semantics describes is thought of here as a relationship between EL sentences. All of which is in marked contrast to the other approach to language, which might be called the psycho-computational tradition, which is primarily concerned with the internal representation.

This more cognitive perspective treats an EL utterance not as having a content, but as producing a content in the mind of the hearer. The meaning, here, is thought of as something to be extracted from the utterance of the sentence in a context. This tradition concentrates on the processes by which external language is understood and the endproduct of those processes, so that EL syntax is thought of as playing a role in the facilitation of comprehension rather than reflecting an objective structure. This tradition is often also concerned with the role of IC in other, nonlinguistic, aspects of cognition. Inference is considered to a relation between propositional structures in IC, so the proper place for a semantics, on this view, is to provide an account of meaning for the end-product of the process of linguistic comprehension, i.e. the internal code, rather than the external communication language<sup>4</sup>. AI work on knowledge representation, for example, has a central concern with the semantics of KR formalisms. In this tradition there is no particular methodological requirement to

<sup>3</sup>For example, Kamp (1990, p. 32): "At present there isn't much that we can say with certainty about the way in which the human mind represents and processes information ... there is little hope that this situation will significantly improve in the foreseeable future... So theorizing about these matters...is something one had better stay away from. ... A theory of attitude reports ought to be independent of any specific assumptions about the organization of mental states and the mechanisms which transform them."

<sup>4</sup>On this view, the grammatical structure of EL might emerge simply as a by-product of the processes of language production and comprehension.

focus on the grammar of EL, although it is often of considerable practical importance<sup>5</sup>.

There is a further difference between the two traditions in what might be called theoretical attitude. Semantic linguistics seeks a theory of the structure of language which is as universal and simple as possible. This leads to a preference for sparse contextual structures containing just sufficient structure to work, and for subsuming as many linguistic phenomena as possible into a common framework. The methodological pressures of computational modeling lead in rather different directions, since the process of comprehension may involve information from many sources, and use this information in ways that depend in part on the source. These distinctions, and the complexities they introduce, can seem like unnecessary clutter to someone working in the first tradition.

#### CODES AREN'T CONTEXTUAL

The grammatical form of an EL utterance directly encodes only what linguistics calls the 'character', i.e. that part of the meaning that can be understood from the sentence in isolation. The character of the sentence "He saw him" is, roughly, that two male creatures exist and one saw the other at some time in the past. A full comprehension of that utterance in a communicative context would use the context to determine the intended referents of 'he' and 'him', resulting in a content which is better

<sup>5</sup>This sketch of these rival traditions deliberately emphasizes their differences rather than their similarities because the differences are the chief barriers to communication, but there are many points of contact between them, especially in more recent work. Linguistic semantics, like linguistics in general, is often motivated by a perceived psychological relevance of the structures it hypothesizes (but rarely accepts the experimental discipline required in psychology or the attention to implementation detail necessary in AI), and computational psycholinguistics sometimes feels itself to have relevance for grammatical theories in linguistics (but rarely pays attention to the syntactic details that interest linguists).

expressed as a ground assertion with no quantifiers, explicitly naming the people being described. This decontextualized content is what IC must be able to encode; it is supposed to represent the final output of the comprehension process, where the full content is represented and stored for later use. The task of the IC is to represent this full content, and provide a vehicle for connecting it to other mental processes. This means that while the process of comprehension uses contextual clues to discover the full content of the message, the IC itself cannot rely on contextual clues to provide meaning in the way that our external, communicative language does, precisely because its role is to encode the content which results from deciphering those clues, and to preserve this content and make it available to subsequent mental operations (in AI, typically, thought of as inferences) after the relevant context is no longer available. The IC should be able to represent information which is completely decontextual in the sense that linguistic semantics uses 'contextual'. It cannot use indexical or anaphoric devices since it must be capable of representing the result of resolving indexicals and anaphoras; it should have no lexical ambiguity since its role is to express the result of lexical disambiguation (Woods 1986); and so on for any contextually sensitive aspect of EL meanings. This has been the standard ambition and usual assumption in the second tradition (although full *lexical* disambiguation has been questioned, as discussed below.)

This point sometimes seems elementary and obvious to those working in this tradition, but incomprehensible or ridiculous to linguistic semanticists. To illustrate it, consider the consequences of allowing the IC to contain indexicals such as 'IC-now.' If the comprehension process encodes the meaning of the present as 'IC-now,' then the *content* of a sentence meaning would also be indexical: the meaning of the EL sentence "It's raining," spoken and understood on Monday as referring to the time of utterance, would be represented using 'IC-now,' and therefore would mean that it is raining on Tuesday when accessed on the following

day. Clearly, part of the comprehension process for such indexicals involves using the context of the utterance to find their non-indexical meaning and then encoding that decontextualized meaning in the internal representational code.<sup>6</sup>

Some care is needed here. While the result of comprehension must be fully decontextualised, the process of language comprehension itself might make use of internal codes which are contextually sensitive<sup>7</sup>. This said, however, the differences between the two traditions are often evident. In the linguistic tradition much recent work has been concerned with the development of under-specified logics intended to encode partial utterance meaning. It is easy to find examples where this lack of contextual precision extends well beyond a single sentence. ("He saw him yesterday." "Who?" "Bob." "Oh hell, what can we do?" "No, *Harry* saw *Bob*. Bob didn't see him, thank God.") These pose particularly acute difficulties for a semantic theory which must define meanings attached to a syntax whose largest unit is the sentence.<sup>8</sup> In AI work, the results of accumulating (and revising) information about the intended meaning are more naturally modeled simply by building progressively more detailed hypothetical meanings in an IC which need have no special connection to the EL syntax. For example, the pre-contextual meaning of "He saw him" might be represented (oversimplifying somewhat) as  $(exists\ x\ y)(Male\ x) \ \& \ (Male\ y) \ \& \ (See\ x\ y)$  thereby encoding the *character* of the EL sentence as the *content* of a fragment of IC; but this

<sup>6</sup> This decontextualized meaning may not be a calendar date, but it must be something which when accessed later will then still refer to the time of speaking rather than the time of access.

<sup>7</sup> I am grateful to an anonymous reviewer of an earlier draft for this observation.

<sup>8</sup> A methodological concern with the sentence as a basic unit, and hence the importance of context effects which transcend sentence boundaries, is characteristic of the linguistic tradition.

content needs no special logics for its expression.

#### PHYSICAL AND LINGUISTIC CONTEXTS

Consider two people in a garden, where one says: "Look at those roses. Aren't they beautiful?" Clearly, 'they' refers to the roses mentioned in the first sentence, and these roses are, in some sense, part of the context which determines the content of the second sentence. However, there are two ways to understand what this means. In one, the context is taken to be the actual physical surroundings of the conversants: the garden itself (or perhaps part of it). Call this a *physical* context. In the other, the context is provided by the words in the first sentence; we might call this the *linguistic* context. If we consider the two sentences as a narrative, 'they' is anaphoric; but if we think of the physical context, 'they' can be regarded as an indexical: it would be just as appropriate, and convey exactly the same meaning, for the speaker to indicate the roses by a gesture or gaze and say: "Aren't they beautiful?", without any linguistic introduction of the topic of flowers.<sup>9</sup>

This contrast between physical and topic contexts is often blurred, or deliberately ignored, in computational linguistics, which uses the term 'common ground' (Clark & Carlson 1981) of a conversation to mean the objects and topics which have been somehow introduced into the range of attention of the participants, plus their mutual beliefs about these objects. Things can be in the common ground either by having been previously mentioned or by being part of the physical surroundings, or indeed by any other means. Linguistic semantics has no need to distinguish these: they play the same role in the subsequent interpretation of sentences by providing a set of possible referents for

<sup>9</sup> Readers with the grammatical sensitivity of many linguists may regard this as ungrammatical, preferring 'Aren't *those* beautiful?' precisely because it has a lexical indicator of the indexical. Nevertheless, the first sentence would be comprehensible in that contextual setting, which is what the second tradition is most interested in.

pronouns, and the semantic constraints arising from shared beliefs apply in either case. One theory neatly subsumes both, so they can be identified.

However, from the cognitive perspective there is an important difference between physical and linguistic contexts. The former, unlike the latter, are liable to change as time passes for nonlinguistic reasons, and once in the past they cannot be recovered. A conversation can return to an earlier topic, using a stack of common ground representations (or perhaps a single common ground itself having a stack-like organization) which can be stored in memory; but when people are speaking about some ongoing process in their physical surroundings, the past cannot be resuscitated for further indexical reference, so the distinction between the anaphoric and indexical usages illustrated above may be crucial. Consider for example the question, "Do you see that?", with no preliminary introduction. The content of the question can be successfully determined only if the referent is visible in the physical environment. If it was a shooting star, the hearer who fails to see it at the time has no way to compute the meaning later.<sup>10</sup> Even more extreme examples of the difference are provided by short command phrases such as "Look out!" or "Stop!", which have only a trivial linguistic structure and can play no role in a narrative, but may convey vitally important information about the physical context.

(It may be objected that whether or not something is physically present, it can only be thought of as part of the mutual context if it has been somehow introduced into the common ground<sup>11</sup>, but this seems to not be true for physical contexts. While narratives

<sup>10</sup>Subsequent conversation may introduce the topic, but the information available to the hearer from being subsequently told that a shooting star had been present is quite different from that obtained from seeing it, and the comprehension process at the time of hearing the question is quite different.

<sup>11</sup>I am grateful to a reviewer for making this suggestion, even though it is wrong.

and third-person accounts of conversations usually take care to first introduce a topic and subsequently refer to it, EL is often used in a natural conversation to comment on something in the physical context which has not been discussed previously, relying on the listener's ability to discover the intended referent from clues in the immediate environment during the comprehension process itself. As well as warning shouts, this is illustrated by examples like: "Have you noticed the Chinese vases over there?", where part of the intended meaning is to precisely direct the hearer's attention to a part of the physical context which is not yet in the common ground.)

From the AI perspective, the key difference between physical and linguistic contexts is that the information in them must be accessed by different mechanisms, which must be sensitive to the different natural structures the two kinds of context have. Linguistic contexts can be represented in memory, stored and accessed later, corresponding to changes in conversational topic. These often seem to obey a stack-like organization, where it is natural to return to the previous topic, whatever it was. In contrast, the meaning of many indexicals can be determined only by perception, and requires close attention to the fleeting nature of physical situations, ordered by the relentless passage of time; but it does not require any internal record of what has been said previously. Consider a transcription of a conversation between two people walking in a garden, or a narrative description of it. To follow this the reader needs to construct a representation of what is going on in the situation being described and modify this representation as the narrative proceeds, to fix the referents of the words spoken by the conversants. However, *they* needed no such representation, since they could see the garden itself: they were already in the ongoing dynamic situation which the reader must somehow model while reading their words. Different plants came into their view simply as a result of their walking. This is particularly clear in a third-person narrative, which can contain explicit information given to the reader by the

author: “ ‘Aren't they lovely?’, said Fanny as she looked happily at the roses”. A conversation has no author, so such information about Fanny's gazing and state of happiness could be obtained only by looking at Fanny, if one were actually listening to her (or by being Fanny.) One might characterize the physical context of a conversation as that part of the common ground in which the conversants are situated, and the linguistic context as the part that is situated in them.

This distinction between physical contexts, which have a linear temporal ordering and provide referents for indexical terms, and linguistic contexts, which have a stack-like nested structure and provide referents for anaphoric expressions, is familiar (see for example the distinction between 'deep' and 'surface' anaphora in Hankamer & Sag 1976). It is less widely appreciated, however, that in one intellectual tradition it is natural to introduce a general category which subsumes them both, while in the other it is natural, and indeed may be essential, to keep them separate.

#### CONCEPTUAL CONTEXTS

In both cases so far we are talking of a context for an utterance or a sentence. A different sense of context means, roughly, the set of background assumptions needed to fix the meaning of an ambiguous word or concept. The different senses of 'bank' in English, for example, include the side of a river, an institution for safely storing cash, and a tilting action of an aircraft. There is much experimental evidence that when we hear a sentence the normal process of comprehension considers all these possibilities and selects the appropriate one in about half a second, presumably somehow using information from the linguistic and physical contexts.

However, it is not exactly clear what constitutes a single 'sense' of a word in EL. In the linguistic tradition, all that is necessary is that the context-structures involved in the semantic theory provide enough information to support any linguistically relevant distinctions, most notably to keep track of

pronoun bindings and distinguish between different syntactic classes. However, if one wishes to encode enough of the content of the meaning of a word in the IC to support a reasonably rich level of subsequent cognitive inference (that is, if the IC is going to be of any actual subsequent use) then things get more difficult. Attempts to formalize such multiple senses often reveal a greater degree of polysemy than is suggested by ordinary lexical classification. For example, even if we restrict ourselves to financial banks, a bank may be a legal institution, a corporation, a legal agent, a building, a collection of buildings, the interior of a building, or even such things as an architectural style. A building may be a physical object which encloses a large space, or it may be that space itself. As one tries to capture these meanings by writing 'axioms' which would be true of them, the meanings tend to separate into finer and finer shades. Is the fitted carpet in the foyer of a bank-building *inside* it or *part* of it? Is a room in a building a *part* of the building or is it an inhabitable space *enclosed* by the building? It has proven very difficult to give convincing logical theories which adequately capture word meanings with sufficient precision to support useful subsequent inferences: there seems to be an "explosion of meaning" (Guha 1990) which taxes our representational abilities. It also seems to be beyond the resolving ability of most linguistic accounts of meaning, which generally assume a much coarser level of lexical ambiguity than that which AI seems to require. Most linguistic semantics would give 'bank' three or four meanings rather than thirty or forty, especially when the distinctions between the meanings seem to have no linguistic significance. This explosion only becomes a problem for the processes that must use the IC to draw subsequent conclusions.

Such problems have led some to conclude that a deductive framework is fundamentally inappropriate as a way to express content; others however (Guha 1990, Singh et al 1995)) have partially abandoned the goal of total disambiguation every utterance, and developed contextual logics to give a coherent account of deduction

in the presence of polysemy. Here, a context is thought of as a kind of semantic index to enable the same logical symbol to be used with several different, but related, meanings. Rather than having to distinguish *bank-single-building*, *bank-building-interior*, etc., these logics would instead allow a single term *bank-building* which can be reinterpreted in various 'contexts' to have different shades of meaning. The role of context here is fundamentally conceptual: to disambiguate the vocabulary sufficiently to enable useful inferences to be made, so I will call these *conceptual* contexts. Notice that conceptual contexts do not eliminate the fine distinctions, but have been proposed rather as a way to keep them properly organized. These contexts are part of the representational IC, not outside it helping to define its meaning.

Physical and topic contexts are found in nature, but conceptual contexts are a formal device. It is natural therefore to ask whether the latter can be used to model or describe the former. If so, the process of meaning resolution in NL comprehension might be performed by inference in a suitable contextual logic (Buvac 1996a explores this possibility.). One way to judge the usefulness of this proposal is to ask what the natural structure of conceptual contexts must be. Just as the fact that we live in time forces physical contexts into a linear ordering, and our ability to return to a previous topic requires linguistic contexts to have a stack-like organization, we might expect that conceptual contexts must fit into some kind of overall structure in order to provide conceptual distinctions. It is not so easy to describe the intended structure of conceptual contexts, but it seems unlikely to be similar to those already considered.

Both physical and topic contexts have a fairly large scope, typically extending over several sentences. The 'common ground' is supposed to include potential referents for all pronouns in a conversation and all the shared beliefs which might be relevant to pronoun resolution, for example. Moreover, they are dynamic, changing as a conversation proceeds, and one of the roles of

EL is to move topics in and out of the narrative context. None of this is true of conceptual contexts. These change, if at all, at a much slower rate: the mental lexicon is largely established in childhood while a language is being learned. Also, the scope of a conceptual context can be very narrow indeed. Quite natural sentences can use many different senses of a single word. For example, each token of the word 'in' in the following sentence has a slightly different sense: "The pen in the tray in the top drawer of the desk in my office was sitting in water a while ago." These all refer to spatial inclusion in one way or another, but the ways are all different and have different deductive properties, so must be somehow distinguishable in the IC<sup>12</sup>. If conceptual contexts are how such distinctions are recorded, then they must cluster to sentence meanings like barnacles on a rock, rather than obeying any straightforward linear or hierarchical organization.

One might object that while their role in disambiguation is local and untidy, the contexts themselves have a natural structure, perhaps reflecting a concept hierarchy or obeying rules of prototypicality. But even here it is very difficult to distinguish any simple structure. All the straightforward early ideas of concept structure (for example, as feature lists) have proven to be inadequate, and the best that psychology can now offer seems to be little more than the observation that different word senses might be organized in families, distinguishable by their satisfying different logical theories. Considered as sets of logical sentences, these 'microtheories' overlap in complex ways. Some of them will contain others, and some groups of theories will have some, but not all, assumptions in common. However, there is no reason whatever to expect them to be

<sup>12</sup>For example, some are transitive, others not. To be in a container in a place implies being in that place; to be an object in a bowl in liquid does not imply being in the liquid, but to be an object in liquid in a bowl does imply being in the bowl. Herskowitz (1986) documents the polysemic intricacies arising from the use of spatial prepositions in English.

temporally ordered. They have nothing particularly to do with time or narrative sequence. And although (like any set of subsets) they can be partially ordered by inclusion, this partial order would not seem to be at all closely related to the topic structure of narrative contexts, for the reasons just discussed: a single line of reasoning, or even a single assertion, will often involve concepts defined in many different conceptual contexts.<sup>13</sup>

## CONTEXT LOGICS AND DEDUCTIVE CONTEXTS

Several authors have proposed adapting first-order predicate logic (the *lingua franca* of knowledge representation in AI) to include some contextual sensitivity. Guha (1991), McCarthy (1993), Buvac (1996b) and others are seeking a general-purpose logic of contexts in which a logical sentence can be asserted only relative to a context, and inference rules are provided for 'moving' between contexts. The relation between a context and a sentence is expressed by a special sentence of the form *ist(k, phi)* (read as 'phi is true in k'). Of course, like any other sentence, this too can be asserted only in a context; but this construction allows some contexts to mention others, and make explicit assertions about what is implicit in assertions made in them. Contexts may have local assumptions assumed to be true within them (in this respect they closely resemble the familiar conditional proof subderivation constructions used in natural deduction logics) and they may also reinterpret some of the vocabulary, so that for example a binary relation may be replaced by a unary property in a subcontext where some parameters are fixed. We might call this kind of context a *deductive* context:

<sup>13</sup>Even here there is scope for computational variation. One view of microtheories imagines them as mostly used for handling exceptional cases, giving a picture of a large common context like a flat savanna, with relatively sparse context/sub-context trees. A very different picture sees almost every assertion as made relative to a microtheory, the total forming a dense structure more like a rain forest or a mangrove swamp. (Guha 1990).

it contains a set of assumptions and specifies a vocabulary which allows inferences to proceed without being hampered by conditional qualifications; a way, as it were, of packaging a set of assumptions and temporary changes to ones vocabulary, like a new suit of axiomatic clothing. McCarthy gives the plausible example of moving from a context in which relations are temporally sensitive into one representing a particular time, where the temporal parameters are considered constant and can therefore be ignored, thus presumably making inferences simpler. McCarthy suggests that counterfactual and fictional reasoning can be expressed in this way, so that one might have a context *sherlock-holmes-stories* within which the fictional characters are assumed to exist, so that

$$\text{ist}(\text{sherlock-holmes-stories},$$

$$(\text{exist } t)(\text{Past}(t) \ \&$$

$$\text{Resides}(\text{Watson}, \text{India}, t))$$

$$)$$

is simply true, i.e. true in our context, even though there is in fact no such person as Watson.

In McCarthy's original account, the logic makes no *a priori* commitments to what connections might be possible between the vocabularies and assumptions of one context and another. The context logic developed by Buvac (1996; Buvac, Buvac & Mason 1995) has a pair of rules which intuitively are used to enter and leave a subcontext from a larger context. The entrance rule translates an assertion about a contextual truth into a simple truth stated within that context:

(Entrance) *From ist(k, phi) infer k: phi*

Since all assertions are made relative to a context, there is here assumed to be an unspoken but implicit global context relative to which any simple assertion is being made; this applies to any inference rule, in fact. The exit rule refers explicitly to this outer context:

(Exit) *From k: phi infer k': ist(k, phi)*

Since the context name k' used here is free, this might be interpreted as meaning that from the sentence *phi* derived in context k,



we can infer that  $ist(k, phi)$  is true in any context; but this would render the logic vacuous, since we could use it to transfer any assertion from any context to any other simply by entering a subcontext from the first and then exiting to the second. McCarthy and Buvac talk of using this rule and its converse to enter a context, perform some inferences, and then exit *back to the original* context. The intended usage is clearly that  $k$  in this rule is a 'supercontext' of  $k$ , and that all context invocations form a nested recursive structure rather like that of subderivations in ND proofs, or the dynamic stacks which are used to make operational sense of recursive subroutine calls in programming languages.<sup>14</sup> (More recent accounts of these logics make this stack structure explicit by labeling with sequences of context names. Appendix A gives an alternative natural-deduction style formalization.)

This stack-like organization suggests that these logics may be useful for describing linguistic contexts, but not for conceptual or physical contexts. Consider for example describing the change in physical contexts relevant to a conversation during a walk through a garden. As the conversants turn a corner, they are confronted by a rose-bed which was previously invisible. This presumably requires a context change, and so we must either enter a new context or exit to an older one, as these are the only kinds of context shift available. But since this situation is new, exiting will be of no use, so entering a new context is the only option. As the walk proceeds, these contexts will

<sup>14</sup> The 'universal' interpretation seems analogous to the rule of necessitation in modal logics: *From p infer Np* which can be understood as saying that a tautology must be true in any possible world. However, this rule is valid only in a simple logic with no mechanisms for introducing temporary hypotheses, since then any sentence in a correct derivation must indeed be a logical theorem and hence necessarily true (if the logic is sound). In context logics, however, the inferences performed in the inner context and extracted to the outer one are typically intended to be correct only locally to that context.

accumulate on the context stack, like ghosts of the way the world was in the past. There is no way (unless the conversants retrace their steps) to get rid of these earlier contexts; the exit rule simply takes us back to the immediately preceding context (now in the past), and the logic provides no mechanism for moving to any other set of assumptions. For conceptual contexts the picture seems even less appropriate, since a cluster of different but overlapping conceptual contexts may be needed to distinguish the several occurrences of a single relation symbol in a single assertion.

Buvac (1996) applies this framework of contextual logic to analyze dialog structure. To this, however, it is necessary to assume that the meanings of the utterances have already been processed into the form of logical sentences. This seems an odd assumption, since that comprehension process is itself contextually located. To be sure, one might expect that the resulting IC would refer to aspects of the context in which the EL was originally interpreted. The IC resulting from the comprehension of a narrative might refer to the sequence of events the narrative describes, for example. But this description need not itself be contextually located: in fact, the very process of comprehension would normally be expected to eliminate that contextual sensitivity, transforming a narrative text like "John got up. He had breakfast. He went out..." into a representation with a meaning expressed more accurately by: "At some time  $t$  John got up, and at time  $t+1$  he had breakfast, and at time  $t+2$  he went out, and...". But this is simply a conjunction, where the narrative time order is made explicit in the content of the IC itself, and does not require any special deductive machinery. (It is instructive to contrast this with the treatment in (ter Meulen 1995), where the narrative structure seems to be intimately involved in the deductive process. terMeulen however is clearly in the linguistic tradition; she is talking of inferences between English sentences, not logical sentences, and the contextual narrative machinery is closely connected with such linguistic delicacies as past and progressive tense inflexions, etc. The logics

do not have these syntactic subtleties, so their meaning must have been already decoded and made explicit in the logical sentences (or lost altogether); but this is only possible - as terMeulen's work clearly illustrates - by taking the relevant narrative contexts into account.

## CONCLUSION

The four senses of 'context' distinguished here seem to clearly be distinct notions: the physical setting of a conversation; the topic of a conversation or narrative; indices for recording fine conceptual distinctions; and finally, temporary packages of deductive assumptions. There may well be other kinds of 'context'. How plausible it is to treat some or all of them as similar, so that a single account can be found which subsumes several of them, seems to depend on ones theoretical assumptions and background. The linguistic tradition typically merges physical and linguistic concepts into a common ground and ignores conceptual and deductive contexts. The computational tradition is more sensitive to the distinctions, but if one thinks of deduction as a universal cognitive-modeling mechanism then it would be natural to hope that deductive contexts might subsume all the others.

A useful general-purpose device is to ask, for each kind of context, what their natural structure seems to be. What purely mathematical structure would describe the significant relations on the set of all such contexts? This question seems to have different answers for the four kinds of context considered here, suggesting that there may be no very useful single notion of 'context', but that like most other English words, its meaning can only be properly understood when the context is made clear.

## REFERENCES

Buvac, S. Ambiguity via Formal Theory of Context, in *Semantic Ambiguity and Underspecification*, ed. Peters & van Deemter, CSLI lecture note , (1996a)

Buvac, S. *Quantificational logic of context*. Proc. 14th Int. Joint Conf. on AI, (1996b)

Buvac, S., Buvac, V. & Mason, I.A., *Metamathematics of Contexts*. *Fundamenta Informaticae* 23(3) (1995)

Clark, H.H. & Carlson, T.B. Context for Comprehension, in *Attention and Performance IX*, ed. Long & Baddeley, pp 313-330. Erlbaum, (1981)

Grosz, B.J. & Sidner, C.L. Attention, Intention and the structure of discourse. *Computational Linguistics* 12: 175-204 (1986)

Guha, R. V. *Contexts: A formalization and some Applications*. Doctoral Dissertation, Stanford University (1991)

Guha, R. V. *Micro-Theories and Contexts in CYC*, Part 1: basic issues. MCC technical report ACTR-CYC-129-90 (1990)

Hankamer, J. & Sag, I.A., *Deep and Surface Anaphora*. *Linguistic Inquiry* 7:391-426 (1976)

Herskowitz, A. (1986) *Language and spatial cognition* Cambridge University Press

Kamp, H. Prolegomena to a Structural theory of Belief and other Attitudes, in *Propositional Attitudes*, ed. Anderson & Owens, CSLI Lecture Notes # 20 (1990)

McCarthy, J. Notes on Formalizing Context. Proc. Thirteenth International Joint Conferences on Artificial Intelligence, (1993)

McCarthy, J & Buvac, S. Formalizing Context, Technical note STAN-CS-TN-94-13, Stanford University (1994)

ter Meulen, A. *Content in Context*. MIT Press (1995)

Singh, N., Tawakol, O. & Genesereth, M. A name-space Context Graph for Multi-Context, Multi-Agent Systems. Working Notes of AAAI Fall symposia, (1995)

Woods, W. (1985) *Whats in a link?* in Brachman & Levesque (1985) *Readings in Knowledge Representation*. Morgan Kaufmann, 1985

## APPENDIX: A NATURAL-DEDUCTION LOGIC OF CONTEXTS

Consider propositional logic extended recursively to include sentences of the form  $ist(k, phi)$  where  $phi$  is a sentence and  $k$  is a constant symbol from a set  $C$  of *context names* disjoint from the rest of the vocabulary. Conjunction and disjunction are understood to be symmetric. Define a *derivation* to be a sequence of *items*, where an item is either a sentence or a derivation labeled with a context name.

We write  $k:[a...b]$  to mean any derivation labeled with  $k$  whose first and last items are  $a$  and  $b$  respectively, and  $k:[...a...]$  to mean any derivation

containing  $a$ . The following rules apply to any items which occur in derivations with the same label, and the conclusion is added to the end of one of those derivations.

-----  
*and:*  
 from  $(a \ \& \ b)$   
 infer  $a$

from  $a, b$   
 infer  $(a \ \& \ b)$

-----  
*or:*  
 from  $(a \ or \ b), k:[a...c], k':[b...c]$   
 infer  $c$

from  $a$   
 infer  $(a \ or \ b)$

-----  
*implies:*  
 from  $a, (a \ implies \ b)$   
 infer  $b$

from  $k:[a...b]$   
 infer  $(a \ implies \ b)$

-----  
*ist:*  
 from  $k:[...a...]$   
 infer  $ist(k, a)$

from  $ist(k, a)$   
 infer  $k:[a]$

-----  
*not:*  
 from  $k:[a...(b \ \& \ not \ b)]$   
 infer  $not \ a$

There are also two special 'structural' rules:

-----  
*hyp:*  
 infer  $k:[a]$   
 if  $k$  does not occur elsewhere in the proof.

-----  
*import:*  
 $k:[...a...k':[...]]$   
 can be transformed to  
 $k:[...a...k':[...a]...]$

This version of contextual logic restricts the contextual sensitivity to making temporary assumptions (although it allows these to be re-used in other co-labeled sub-derivations without needing to be repeated there). More elaborate kinds of contextual restriction would require a more elaborate syntax, and this use of contexts, particularly in the negation and implication-introduction rules, would then be restricted to the simple case.

It is straightforward to extend the usual completeness, compactness, etc. proofs to this logic as it is equivalent to a standard ND formulation of propositional logic. (The quantifiers pose no special problem but are omitted here to save space.) The point of exhibiting it here is to emphasize that these deductive contexts play an essentially proof-theoretic role, and need have no special connection to the process of thought, even if that process is expected to produce proofs which are valid in this logic. A theorem-prover for this logic need not be closely related to the contextual

inclusion relation, but could proceed, for example, by backward chaining. The *ist* rules need not be interpreted as entry and exit processes, and doing so imposes an unnecessary extra burden on the intended meaning.