The Sociality of Context

Bruce Edmonds

Centre for Policy Modelling,
Manchester Metropolitan Univeristy
bruce@edmonds.name

Abstract. It is well established that many human abilities are context-dependent, including: language, preference judgement, memory, reasoning, learning and perception. This is usually taken as a negative – that there will be limits on our understanding and modelling of these abilities. However, what is not always appreciated is that context-dependency can be a powerful tool in social coordination and communication. This paper pulls together several theories about the cognition of context, and presents a computational model of context-dependency. It then sketches its role in social communication, coordination and embedding. It looks at some of the approaches to dealing with context in the computer science and social science literature and concludes that none of these squarely faces the problem of context dependency. This points towards a substantial gap in the research and hence a future programme.

About Context and Context-Dependency

Context pervades the human social and cognitive realms but due to its very nature it is often unnoticed or left implicit.

The implicit nature of context means that attempts to label it are often "over-loaded", with the result that the word "context" seems to have a variety of related, but distinct, meanings. Here I do not want to enter the debate concerning the "right" meaning of this word, since that has turned out to be a fairly fruitless enterprise, but I do want to make clear the way in which I am approaching the subject. Unfortunately, "Context" is used in many different senses and has many different analyses. It is somewhat of a "dustbin" concept, in that if a theory or idea does not work the reason may be assigned to "the context". Thus to many (e.g. linguists) context is a subject that is to be avoided due to its difficulty. "Context" is closely related to (but not identical to) a number of other concepts, including: tacit knowledge (Polanyi 1966), the frame problem (McCarthy and Hayes 1969), framing in psychology (Goffman 1974), and the "situation" (Barwise and Perry 1983).

Situational Context

The situation context is the particular situation where some events or other described phenomena takes place. This could include the time and location, but could include all that is the case about that situation, including: who was there, the knowledge of

those people, the history of the place and all the objects present. In this sense the context is indefinitely extensive, it notionally includes all the circumstances in which an event or utterance occurs.

Such a context may be able to be specified adequately (if rather uninformatively) by giving the time and place of the events, but the relevant details might not be effectively retrievable from this. For example, the fact "I was reminiscing about our summer holiday" might well not be detectable from the time and place except by the person doing the reminiscing. Thus when talking about the situational context it is almost universal to abstract from this to what is relevant about that context, or what might be commonly understood (and hence safely not described but left implicit). Thus the phrase "the context" (as in the question "what was the context?") may *mean* "those factors that are relevant to understand this particular occurrence" even though it may *refer* to the situational context in general. Thus to understand what someone is saying to you, you might ask "what was the context?" and get a description of the relevant features of the circumstances, e.g. "I was on the train home".

Linguistic Context

Whilst the situational context could include anything, at least in theory, the linguistic context is composed of the words that surround an utterance or phrase. This typically indicates the words that precede or frame the target of understanding, but could also include the common knowledge that could be reasonably be expected to be known by the listener/reader, e.g. elements of the relevant culture. Sometimes this is taken to be the same as all that which is necessary to understand some natural language.

Historically this has been the last resort of the linguist in attempting to pin down the meaning of an utterance – what one appeals to if there seems to be no detectable foreground features to explain its meaning. However recently more positive attention has be focused on context in linguistics. For example, Peter Gardenförs (1997) has said (pragmatics being close to contextual considerations in linguistics):

> *Action is primary, pragmatics consists of the rules for linguistic actions, semantics is conventionalised pragmatics and syntax adds markers to help disambiguation (when context does not suffice).*

Clearly the linguistic context could refer to almost any of the language or culture that surrounds an utterance, and hence is not something that can be captured in its entirety. Often context is thought of as linguistic context because the social interactions that are being considered are composed of linguistic communications.

Cognitive Context and Framing

Clearly many aspects of human cognition are context-dependent, including: visual perception, choice making, memory, reasoning and emotion (Tomasello 1999, Kokinov and Grinberg 2001). A lot of recall, learning and inference is only done *with respect to* a recognised kind of situation. That is, some knowledge is acquired in a particular situation and then made available in similar situations. This abstraction of a situation in the brain – the recognised kind of a situation in which packages of

knowledge etc. are relevant – is sometimes called the "cognitive context" (Hayes 1995). This is the cognitive correlate of the situational or linguistic context. Such cognitive contexts could be identified using a description of the kind of situational context that invokes them or else by the set of all the knowledge, norms, expectations, habits etc. that are immediately accessible once recognised. Humans seem to have an innate ability to recognise the cognitive contexts of others (Tomasello 1999).

This idea of cognitive context relates to the idea of framing in psychology as introduced by Goffman (1974), where frames are conceptualised as "schemata of interpretation" that allow individuals "to locate, perceive, identify, and label" experiences. In this way it can be said go make explicit thought possible by giving meaning, organizing these experiences, and guiding actions. This implies that framing is what happens to thought when a cognitive context is adopted. Indeed (Elliott & Hayward 1998) state the following.

> *"Frames are associated with, or even equivalent to, the numerous social and cultural contexts which define the appropriateness of particular norms of behaviour." (p.234)*

However, as (Entman 1993) points out there is no agreed definition of framing. Kahenman and Tversky (1984) use a clear but simple formulation of framing but this is restricted to the use of different reference points with respect to subjective utility of gains and losses. Shafir, Simonson and Tversky (1993) propose a more general "reason-based" model of choice in which

> *"Different frames, contexts and elicitation procedures highlight different aspects of the options and bring forth the different reasons and considerations that influence decision" (p.34)*

Thus the action of framing can be seen as one effect of considering something within a particular cognitive context – it is the result of cognitive context on opinion and choice. The two concepts are very close but the idea of Cognitive Context is more general – it encompasses other areas, such as how these patterns of salience are acquired and how it affects the acquisition of knowledge as well as its application.

It is essential that different kinds of contexts can be effectively and reliably recognised as cognitive context but this does not mean that they have to be consciously recognisable as distinct cognitive contexts and labelled. For example, they may be unconsciously recognised by all the members of a community but never named; or maybe they their features are distinctive and consciously recognisable but too complex and fuzzy to be completely specified.

Identifying and Talking about Context

One of the difficulties in discussing cognitive context is that they may well not (a) be accessible to us (b) identifiable even if they are accessible or even (c) definable in precise terms even if we can identify them. Thus although, in some way, the brain abstracts its incoming stream of information to some properties of its state that it can later use to recognise and retrieve knowledge that is relevant to the same kind of situational context, there is no reason to suppose that we can safely reify these

properties that would correspond to the cognitive context. Rather we often have to try and deduce what the cognitive contexts are by introspection and other observation.

Despite this, we often talk about contexts as if they were discrete and identifiable "things", however it needs to be understood that for our conscious selves they may not be the case. Thus "the" context is an abstraction of the aspects of those background features that define it, whether or not this is a meaningful or reifiable entity for us. To simplify the discussion I will generally talk about contexts in the sections below as if they are well defined identifiable entities, but the caveats just mentioned need to be always taken into account. This difficulty means that the cognitive context for any situation is often not made explicit or represented – those involved may well not be aware of the cognitive context they are assuming.

The fact that the relevant cognitive context may not be directly accessible to our consciousness does not mean that it is totally immune to being partially identified or uncovered, just that this might be unnoticed, non-obvious, complex, fuzzy and only partially inferable. For example, although we may not be aware of what brought to mind a particular person in a situation, on introspection we might be able to work out that some music brought to mind a past event in which that person figured. Thus we may be able to work out something about what sort of cognitive context is relevant but still not be able to characterise it completely.

A Model of Contextual Cognition

In this section I look at the outlines of a lightweight cognitive model that allows for context-dependent cognition to be implemented within social simulation models. This model integrates Machine-learning type of learning with an AI kind of reasoning via a context-structured memory.

Both learning and reasoning are far more feasible when their scope is restricted to a particular context because this means that only the relevant knowledge needs to be dealt with. However if any significant degree of generality is to be obtained in this manner (McCarthy 1971) then an intelligence must be able to appropriately change this focus as the external context, that is the context we inhabit in (Barwise and Perry 1983), changes. In other words there needs to be some internal correlate of the external context that allows an intelligence to identify which set of beliefs apply. We will call this internal correlate the cognitive context – this is the "internal" approach identified in (Hayes 1997). There are (at least) two tasks necessary for this:

- identifying the appropriate cognitive context from perceptions;
- accessing the appropriate beliefs given the identified cognitive context.

The success of this strategy of assessing the relevance of knowledge via identifiable "contexts" depends upon whether the environment is usefully divided up in such a manner. This is a contingent matter – one can imagine (or devise) environments where this is so and others where it is not. The "pragmatic roots" of context, i.e. why context works, depends upon the underlying pattern of commonalities that occur in an environment or problem domain (Edmonds 1999a). A cognitive context indicates the boundaries of what might be relevant in any situation.

Context serves not only to make it feasible to deal with our knowledge at any one time but also, at a more fundamental level, to make our modelling of the world at all feasible. The efficacy of our limited learning and inference in dealing with our complex world is dependent on the presumption that many of the possible causes or affects of events that are important remain relatively constant (Zadrozky 1997). Otherwise we would need to include all possible causes and affects in our models and decision making processes, which is clearly infeasible. It is the existence of relative constancy of many factors in particular situations that makes our limited modelling ability useful: we can learn a simple model in one circumstance and successfully use it in another circumstance that is sufficiently similar to the first (the same "context").

It is the possibility of the transference of knowledge via fairly simple models from the circumstances where they are learnt to the circumstances in which they are applied that allows the emergence of context. The utility of "context" comes from the possibility of such transference. If this were not feasible then "context", as such, would not arise. For such a transference to be possible the following must hold:

- some of the possible factors relevant to important events are effectively separable;
- a useful distinction can be made between those factors that can be categorized as foreground features and the others (the constant, background features);
- similar background factors are recognizable later on as the same "context";
- the world is regular enough for such models to be learnable;
- the world is regular enough for such learnt models to be useful where such a context can be recognized.

While this transference of learnt models to applicable situations is the basic process, observers and analysts of this process might identify some of these combinations of features that allow recognition and abstract them as "a context". Note that it possible that such an observer will not be able to do this consciously, as it may be obscure, too complex or difficult to analyze into definable cases.

Such a strategy answers those of the "frame problem" (McCarthy and Hayes 1969). Firstly, although the frame problem may be unsolvable in general it is learnable in particular contingent cases. Secondly, the identification of appropriate contexts is not completely accessible to reasoning or crisp definition but is an unreliable, information-rich, and imprecise process. Knowing B in context A, is not expressible as A→B, because the A is not a reified entity that can be reasoned about.

The power of context seems to come from this combination of "fuzzy" and fluid context identity and crisp, relatively simple context "contents". Thus context straddles the fields of Machine Learning and Artificial Intelligence. Machine learning seems to have developed appropriate methods for complex and uncertain pattern recognition suitable for the identification of context. Artificial Intelligence has developed techniques for the manipulation of crisp formal expressions. Context (as conceived here) allows both to be used for different functions in a coherent way.

Combining Context-Dependent Learning and Reasoning

Restricting both reasoning and learning to an appropriate context makes both more feasible. However, as with any other technique, there are a number of difficulties with applying a context-dependent approach to reasoning. *Firstly*:

> *Explicitly specifying a set of knowledge appropriate for a whole set of potential contexts is both time-consuming and labor-intensive.*

Thus with a few honorable exceptions (e.g. CYC (Lenat 1995)), most systems of context-dependent learning or reasoning are only tried out with a few contexts. A possible answer to this (and the one employed here) is to learn the contexts and the context-dependent knowledge. The second is easier than the first; for there are a number of techniques to learn the knowledge associated with contexts.

The learning of the contexts themselves (i.e. how to recognize when a set of beliefs learnt in a previous situation are again applicable) requires a sort of meta-learning. As documented above, there are such techniques in existence. However most of these either require reasonably strong assumptions about the particular nature of the contexts concerned. An exception is (Edmonds 2001b) which describes how contexts can be co-learnt along with the knowledge associated with those contexts. This applies an evolutionary learning algorithm where the knowledge is distributed across a space, where different positions in that space are associated with different set of perceptions or different parts of a problem. This can be clearly understood via the following ecological analogy. If the space can be thought of as a landscape where different parts of the landscape have different properties, and different plants require different properties (some might thrive in marshy land, others sunny and dry etc.). The set of solutions can be seen as varieties of a plant. The different varieties propagate and cross with others in each locality so that, eventually, each variety adapts and, at the same time, spreads across the areas that it is best adapted for. The patches where different sets of varieties thrive define the different ecological niches – corresponding to the different contexts via this analogy.

The ability to learn context allows us to progress beyond the 'loose' loop of:

```
repeat
    learn/update beliefs
    deduce intentions, plans and actions
until finished
```

to a more integrated loop of:

```
repeat
    repeat
        recognise/learn/choose context
        induce/adapt/update beliefs in that context
        deduce predictions/conclusions in that context
    until predictions are consistent
            and actions/plans can be determined
    plan & act
until finished.
```

Such a co-development of cognitive contexts along side their "contents" gives rise to a new problem when the knowledge in these contexts is used to infer predictions and decisions. Thus a *second* problem is this:

> *When some of the contents turn out to be wrong, how can one tell when it is the context that is wrong and when it is the contents that are wrong?*

There is no universal answer to such a question – it will, in general, depend upon the nature of the domain and hence the appropriate contexts in that domain. However there is a heuristic, as follows: if only a few of the elements of knowledge associated with a context are disconfirmed, it is likely that these are wrong (update the set); if many of the elements are dis-confirmed then it is likely that the context is wrong (change it and learn from this).

Thus in the proposed architecture there are four modules: (1) the context identification system; (2) the context-dependent memory; (3) the local learning/induction algorithm; and (4) the inference system, as shown in **Error! Reference source not found.**.

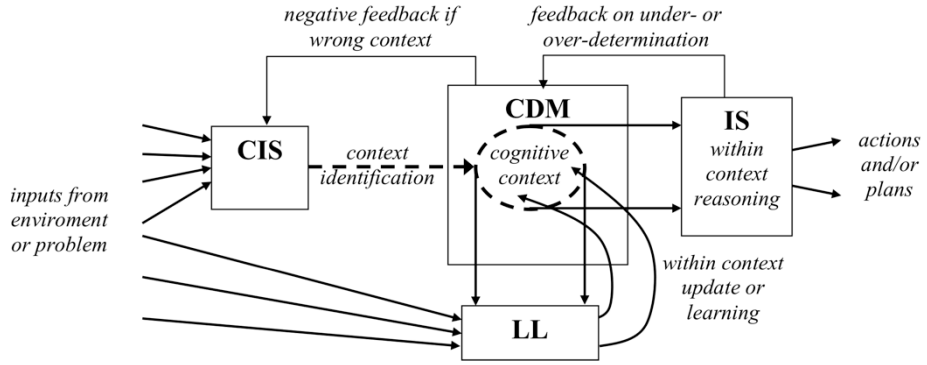

Figure 1. How the context-identification system (CIS), the context-dependent memory (CDM), the local learning algorithm (LL), and Inference system (IS) work together.

The context identification system (CIS) takes a rich range of inputs and learns in a flexible and imprecise way an indication of the context (which it outputs to the memory). The CIS learns as the result of negative feedback when too much of the knowledge in the cognitive context is disconfirmed.

The context-dependent memory (CDM) takes the indication given by the CIS and identifies all those memory items stored within that context. It evaluates the (current) truth of these and if too many are false it returns negative feedback to the CIS, which will identify another context. If a sufficient number of indicated contents are true, then the local learning updates the items within that context. Those items that are (currently) true are passed to the inference system.

The local learning algorithm (LL) performs a local update of the knowledge in the memory. It may include the propagation of successful items towards the focus, but may also include the deletion/correction of items that were false and the possible insertion of new induced/learned.

Finally the planning/inference system (IS) tries to deduce some decisions as to the actions or plans to execute. It could do this in a number of ways, but this could include trying to predict the future states of the world given possible actions and comparing the predictions using its goals.

Two common problems with inference systems that attempt to deduce predictions or decisions from an arbitrary collection of knowledge are under- and over-determination. Under-determination is when there is not enough information to come

to a conclusion or decision that needs to be reached. In other words there may be a key proposition, α, such that neither α nor $\neg\alpha$ can be inferred. Over-determination is when there is contradictory information in the knowledge, i.e. when there is an α such that both α and $\neg\alpha$ can be deduced.

This architecture allows a useful response in these two situations. In the case of under-determination the context can be expanded so that more knowledge can be made available to the IS so that it may make more inferences. In the case of over-determination the context can be reduced so that some of the knowledge can be excluded, the knowledge that is peripheral to the context.

Many non-monotonic logics can be seen as attempts to solve the above problems in a generic way, i.e. without reference to any contingent properties obtained from the particular contexts they are applied in. So, for example, some use 'entrenchment' to determine which extra information can be employed (e.g. oldest information is more reliable (Gärdenfors 1984)), and others allow a variety of default information to be used (e.g. using extra negative knowledge as long as it is consistent (Reiter 1980)). These may work well on occasion, but they cannot exploit any of the relevance relations specific to the particular knowledge and context.

Learning Context

In order for context-dependent reasoning to occur, the context-dependent information (or beliefs) needs to be captured. If the relevant contexts are already known by the designer (and one can recognise when they apply), then either the relevant information can be entered or a context-enhanced learning algorithm can be employed to learn the information with respect to each context. The former case can be onerous because one not only has to enter the relevant facts as well as specifying each fact's domain of application, but one also has to define all the 'lifting-rules' to allow the integration of the context-dependent information. In the later case the context-dependency of the learning means that one needs correspondingly more information within each context for the learning to be complete.

Thus in order for the desired efficiency in terms of context-constrained reasoning to occur (without a laborious entry of information) for each appropriate context, this information (that is both the contexts and the content in the contexts) should be learned by the agent, at least to some extent.

The basic idea is to simultaneously learn the models and the circumstances in which they work best. If there is sufficient regularity in the environment to allow it this will allow some clusters of similar circumstances to be identified and the corresponding models to be induced. However the clustering and induction parts of the algorithm can not work independently; i.e. clusters of like circumstances being identified and then models induced for these clusters. The reason for this is the contexts are identified by those circumstances where particular models work best. These may correspond to a neat (i.e. humanly identifiable) cluster but this is not inevitable – they may be (to the human eye) inextricably intertwined or overlapping. There is not space to describe a possible algorithm for this here (Edmonds 2001b).

The Social Uses of Context

Social Context

Some of the cognitive contexts we have learnt seem to correspond to recognisable kinds of social situation. Examples include: greeting, lecturing, and a political discussion. Once established these seem to be self-perpetuating, in that habits, conventions, norms, terms etc. can be developed by people who recognise the context, but in turn this might mean that the context is more recognisable as an important kind of situation which has its own characteristics. Thus social contexts can be co-constructed over time and passed-on (in terms of experience and social artefacts).

When people are asked to describe the context of an event, they will often do it in social terms. Thus it is that the social context, although it is a special case of situational context is closely linked to the synchronised cognitive context that participants have learnt to associate with situation, because it is often the social aspects that are important in terms of communication and understanding. It is because of the context-dependency of human cognition that when the social context is recognised, experienced participants in that social context will know what set of norms, habits, terms, etc. are associated with it and automatically bring them to bear in their social behaviour and mutual organisation. Thus one of the consequences of the context-dependency of our cognitive capabilities is the prevalence and importance of social context in our understanding of the world.

Some contexts become socially entrenched. That is to say over time a particular kind of situation is supported, perpetuated and marked so that it perseveres and is easily recognisable. This seems to be a co-development of the following processes.

1. The context is made more recognisable by the use of signs, borders, dress, and particular behaviours particular to that kind of situation
2. The meaning and effects of the activities that happen there are valued by the individuals involved and maybe wider society
3. Particular habits, language, behaviours, norms, and tools are developed for that particular context

Thus, for example, a lecture is a socially entrenched context. We often mark it out with a particular location, with its own walls, specially designed seats, technology and layout. Particular kinds of behaviour and norms are in force there, which we train into our children from an early age. It is almost universally recognisable and has come to play a number of important roles in our society.

Such entrenched social contexts will often be explicitly labelled and demarked. They can take on important ceremonial and symbolic meanings within a society. Understanding the significance of such contexts is an important part of being a part of a culture: recognising such entrenched contexts is easy, but knowing the appropriate norms, language, emotions, and habits can take a considerable period of acculturation.

Communication

Although communication used to be thought of as a process of explicit symbol passing, any communication relies upon a previously established commonality or coding structure. How humans communicate is still not entirely clear, but it is becoming increasingly clear that most human language relies heavily on pragmatics and context with the syntax often only used to disambiguate when other factors are not enough (Gärdenfors 1997). Indeed, a lot of communication does not pass any information but is more like an action (Austin 1975) whose use and effect is context-dependent. Whether or not the background knowledge of society, upon which much communication relies, has to be acquired in the painfully slow way humans do, for computational devices to be able to communicate fluently and appropriately will require them to be able to recognise the cognitive context of those they are communicating with (Tomasello 1999), as most six year old children can do.

Social Reasoning and Justification

As Toulmin (1958) recognised most reasoning that is used in a social context does not correspond well to most logical formalisms, but rather to a web of argumentation that justifies (rather than proves) the conclusions from some data given a warrant. The warrant and data can be further argued about producing a web of further supporting and attacking arguments. This process does not stop at any 'natural' ground level of "facts" but rather upon the swath of commonly agreed beliefs. The set of applicable shared beliefs will be dependent upon the context – it may be acceptable to assume that there is a centralised chain of command at sea, but not when at a party. Most socially exhibited reasoning is either to persuade others of a certain action or to explain/justify an action already taken. This kind of reasoning is not logical in the sense of formal logic, but a network of context-dependent argumentation grounded upon the shared knowledge relevant to the kind of situation. To participate in such social discussion as to action computational devices will have to have access to these beliefs – and if they have to do this in a way that is sensitive to changing circumstance, will need to recognise appropriate context and access the appropriate shared belief set according to that context.

Coordination

The social intelligence hypothesis (Kummer et al 1997) suggests that the features of the human brain that have allowed our species to survive are its social abilities. The ability to social coordinate allows us to over come more or avoid our predators and gain access to many different sources of food and energy. However one of these is the ability to develop a "culture" of tools, ceremonies, knowledge, social norms etc. that allow a group of humans to inhabit very different ecological niches (Reader 1990), and thus some groups of human can survive even give unpredictable and extensive catastrophes. However such cultures are not composed of a uniform and general set of knowledge, but overlapping sets that are applicable in different contexts – what is applicable when fishing in quiet waters will probably not be the same as

what holds if the wind gets up to dangerous levels. The ability of humans to co-develop then co-recognise the same contexts and to learn what knowledge is applicable in those various contexts allows for more than shared mechanisms of coordination – it allows for context-sensitive ways of coordination. Thus different ways for people work together can be developed for different kinds of situation. Culture, in its widest, anthropological, sense, does not only consist of a wealth of different beliefs, language, norms, habits, signs, emotions etc. but also a complex of when different sets of these are applicable – and even when the same situation can be considered as alternative kinds of context (e.g. thinking of an invited lecture as an interview for a job). If computational devices are going to be able to participate in such coordination they are going to have to be sensitive to this web of contexts if they are going to truly participate.

Conclusion

Although often treated as a negative – something that delimits our ability to generalise – context is central to human cognition and society. If computational devices are to 'fit' well into our society, they will have to be given some power to recognise and use similar kinds of context in a similar way to us. However, regardless of how we want to develop such devices, a computational understanding of human society will necessitate coming to grips with context and not leaving it to the side lines.

References

Austin, J.L. (1975) How to do Things with Words, Second Edition. Oxford: Oxford University Press.

Barwise, J. and Perry, J. (1983). Situations and Attitudes. Cambridge: MIT Press.

Edmonds, B. (1999a) The Pragmatic Roots of Context. CONTEXT'99, Trento, Italy, September 1999. LNAI, 1688:119-132.

Edmonds, B. (2001b) Learning Appropriate Contexts. In: Akman, V. et. al (eds.) Modelling and Using Context - CONTEXT 2001, Dundee, July, 2001. LNAI, 2116:143-155.

Edmonds, B. (2002) Learning and Exploiting Context in Agents. 1st Conf. on Autonomous Agents and Multiagent Systems (AAMAS), Bologna, Italy. ACM Press, 1231-1238.

Edmonds, B. and Norling, E. (2007) Integrating Learning and Inference in Multi-Agent Systems Using Cognitive Context. In Antunes, L. and Takadama, K. (Eds.) Multi-Agent-Based Simulation VII, 4442:142-155.

Elliott, C. S. & Hayward, D. M. (1998) The exanding definition of framing and its particular impact on economic experimentation. Journal of Socio-Economics, 27(2), 229-243.

Entman, R. M. (1993). Framing: Towards clarification of a fractured paradigm. Journal of Communication, 43, 51-58.

Gärdenfors, P. (1984) Epistemic Importance and Minimal Changes of Belief. Australasian Journal of Philosophy, 62 (2). 136--157.

Gärdenfors, P., (1997) The pragmatic role of modality in natural language. in 20th Wittgenstein Symposium, (Kirchberg am Weshel, Lower Austria, 1997), Wittgenstein Society.

Gilbert, N. (2006) When Does Social Simulation Need Cognitive Models? In Sun, R. (ed.) Cognition and Multi-Agent Interaction: From Cognitive Modeling to Social Simulation. Cambridge University Press, 428-432.

Goffman, E. (1974). Frame Analysis: An essay on the organization of experience. Cambridge: Harvard University Press.

Hayes, P. (1995). Contexts in Context. Context in Knowledge Representation and Natural Language, AAAI Fall Symposium, November 1997, MIT, Cambridge.

Kahneman, D.; Tversky, A. (1984). "Choices, values and frames". American Psychologist 39 (4): 341–350.

Kokinov, B. and Grinberg, M. (2001) Simulating Context Effects in Problem Solving with AMBR. in Akman, V., Bouquet, P., Thomason, R. and Young, R.A. eds. Modelling and Using Context, Springer-Verlag, LNAI 2116:221-234.

Kummer, H., Daston, L., Gigerenzer, G., & Silk, J. (1997). The social intelligence hypothesis. Weingart et. al (eds.), Human by Nature: between biology and the social sciences. Hillsdale, NJ: Lawrence Erlbaum Associates, 157-179.

Lenat 1995 CYC: a large-scale investment in knowledge infrastructure. Communications of the ACM, 38(11):33-38.

McCarthy, J. (1971) Generality in Artificial-Intelligence - Turing Award Lecture. Communications of the ACM, 30 (12). 1030-1035.

McCarthy, J. and Hayes, P.J. (1969). Some philosophical problems from the standpoint of artificial intelligence. Machine Intelligence, 4:463-502.

Polanyi, M. (1966). The Tacit Dimension. Doubleday.

Reader, J. (1990) Man on Earth, Penguin books.

Reiter, R. (1980) A Logic for Default Reasoning. Artif Intell, 13. 81-132.

Shafir, E., Simonson, I, & Tversky, A. (1993). Reason-based choice. Cognition, 49 (1-2), 11-36.

Tomasello, M. 1999 The cultural origins of human cognition. Cambridge, MA: Harvard University Press.

Toulmin, S.E. (1958) The Uses of Argument, Cambridge University Press.

Zadrozny, W. (1997). A Pragmatic Approach to Context. Context in Knowledge Representation and Natural Language, AAAI Fall Symposium, November 1997, MIT, Cambridge.