

HELPFUL ANSWERS TO MODAL AND HYPOTHETICAL QUESTIONS

Anne De Roeck, Richard Ball, Keith Brown, Chris Fox, Marjolein Groefsema, Nadim Obeid, Ray Turner

University of Essex
England

email: deroe@uk.ac.essex (janet)

1.0 ABSTRACT

This paper describes a computational pragmatic model which is geared towards providing helpful answers to modal and hypothetical questions. The work brings together elements from formal semantic theories on modality and question answering, defines a wider, pragmatically flavoured, notion of answerhood based on non-monotonic inference and develops a notion of context, within which aspects of more cognitively oriented theories, such as Relevance Theory, can be accommodated. The model has been implemented. The research was funded by ESRC grant number R000231279.

Keywords: Semantics, Pragmatics

2.0 INTRODUCTION.

Answers people give to questions have two basic properties: they may vary depending on the situation a question is asked in, and, especially if the answer is negative, they aim to be "helpful". The context-sensitivity of answering seems obvious and in no need of further demonstration. What precisely constitutes "helpfulness" is harder to pin down. Modal and hypothetical questions offer an interesting area for investigating "helpfulness". Suppose A and B are going to a party and are discussing how they might travel. Suppose A asks B *Can you drive?* B is correct but perverse to respond *Yes*, if he knows how to drive, has a valid licence but has no car, or if he has a car but he has lent it to someone. A more helpful answer might be *No, because I haven't got a car*. Note here that there is a range of "correct" answers, some of which are *Yes*: for instance *Yes, I know how to drive*, *Yes, I have a licence*, even *Yes, I have a car*, and some of which are *No*, as in *No, I haven't got a car tonight*. The range includes *No, but I can ask to borrow my wife's car*, or even *Yes, if I can get my car back* which establishes a link with hypothetical questions. Note also that, for each of these "correct" answers, we can imagine contexts in which they would be "helpful".

Little is known about the nature of questions and their relationship to appropriate answers, or about how such answers can be computed given some information about what the answerer knows. Some theories, mainly emerging from the Montague tradition (see Groenendijk and Stockhof [1984]), attempt to define "semantic" answerhood (see section 2.2), but fall short when tackling the pragmatic aspect of helpful answers. Other theories [Sperber & Wilson 1986] offer interesting pragmatic insights but their formulation does not allow for a straightforward implementation. Furthermore, the problem of answering modal and hypothetical questions is a compounded one which touches on a host of issues including quantification, intensionality, partiality, belief revision, propositional attitudes, etc.

Our research aimed to draw up a formally specified and computationally feasible pragmatic theory which could accommodate formal semantic views on answerhood as well as intuitive insights into "helpfulness" and its dependence on context (such as offered by Relevance Theory [Sperber and Wilson 1986]). Furthermore, the model is rigorously constrained as it must be tested by an implementation over some knowledge base representing what an agent knows.

This paper is intended as an overview of the computational model. As such it does not provide an in-depth account of all as-

pects of the investigation; in particular, it does not attempt to give a formal account of a basic theory of pragmatics, which is available elsewhere [Ball et al 1990]. Rather, we sketch the background to the problems involved in providing helpful answers to modal and hypothetical questions as a review of the relevant literature and its perceived shortcomings. We will then proceed to outline the intuitions behind our approach to a model of pragmatics and its implementation, and explain how it accommodates helpful answers to modal and hypothetical questions. An example is presented.

3.0 HELPFULNESS, MODALS AND QUESTIONS

Though the problem of helpfully answering modal questions touches on many issues, four particular points need to be addressed.

3.1 Modality.

When looking at the example set out in the introduction, the question arises whether the range of "correct" answers to some extent corresponds to ambiguity (ability/possibility) residing in the modal *can*. Indeed, a large proportion of the literature on modality concerns the view that modals are polysemous, depending on the kind and degree of modality they express (epistemic, deontic, etc.) [Palmer 1979, 1986; Quirk et al. 1985]. Usually, attempts are made to identify modal "primitives" (*ability, permission, etc*) and to analyse modal constructs as ambiguous over several "literal" meanings involving these primitives. Invariably, polysemic reductionist approaches to modal constructs run into problems: given any classification of core types of modality, it is often impossible to determine which reading is involved in any particular example [Coates 1983 versus Walton 1988].

Kratzer [1977] takes another view. She presents a unified analysis of modality which includes the treatment of conditionals (and hence hypotheticals). Modals are unambiguous and modal constructs are analysed as tri-partite structures [see also Partee 1988, Heim 1982], comprising a modal operator, a conversational background, and a proposition. For example, in [Kratzer 1977] the modal *must* in the following sentences:

- (a) All Maori children *must* learn the names of their ancestors.
- (b) The ancestors of the Maoris *must* have arrived from Tahiti.
- (c) If you *must* sneeze, at least use your handkerchief.
- (d) When Kahukura-nui died, the people of Kahungunu said: Rakaipaka *must* be our chief.

is traditionally analysed as (a) 'deontic' *must* indicating duty, (b) 'epistemic' *must* referring to a piece of knowledge or information, (c) 'dispositional' *must*, referring to people's dispositions (e.g. they cannot help sneezing), and (d) 'preferential' *must* referring to preferences and wishes. Kratzer points out that classifications of modals drawn in the polysemous paradigms never adequately cover the data and that new examples are easily found to demonstrate the need for ever more refined categories of modal meaning.

Kratzer wishes to propose a treatment that brings out the common factor in all uses of *must* (and of other modals) and suggests that the burden of differentiation is to be placed on a variation in

context. As such, the meaning of (a)-(d) entertains a relationship with the meaning of (a')-(d') respectively:

- (a) *In view of what their tribal duties are*, all Maori children *must* learn the names of their ancestors.
- (b) *In view of what is known*, the ancestors of the Maoris *must* have arrived from Tahiti.
- (c) If, *in view of what your dispositions are*, you *must* sneeze, at least use your handkerchief.
- (d) When Kahukura-nui died, the people of Kahungunu said: *In view of what is good for us*, Rakaipaka *must* be our chief.

and she defines the semantic interpretation underlying modal constructs as a tripartite structure (applied to (b)):

Sentence: Operator: *Must in view of*
First Argument: *What is known*
Second Argument: *The ancestors of the Maori have arrived from Tahiti.*

The modal is an operator which takes a context and a proposition. The truth conditions for *must*, interpreted as *necessity*, dictate that the modal construct is true if the proposition (the second argument) logically follows from the context (the first argument). A similar approach to *can* (*possibility*) unpacks its truth conditions as true if the context (the first argument) is logically compatible (i.e. does not induce a contradiction) with the proposition (the second argument). Kratzer works within the classical possible world tradition. Conversational backgrounds, modelled as sets of propositions, are usually implicit and linked to the utterance situation, though it is not clear by what mechanism.

Kratzer [1983] proceeds to distinguish between different kinds of conversational backgrounds, depending on the information they contain. She does however experience difficulties when trying to identify different context classes. Indeed, it is as difficult to isolate different conversational backgrounds as it is to pinpoint the various meanings a modal might have.

3.2 Questions and Answers

It is also necessary to present a coherent perspective on questions and answers. Groenendijk and Stockhof [1984] compile an overview of various treatments of questions that populate the field and investigate desiderata for a semantic theory. They argue that interrogatives are entitled to a meaning of their own (and should not be viewed as, say, hidden imperatives) but that their treatment must show some equivalence with that of indirect questions. The meaning of a question is to be related inextricably with its answerhood conditions. Groenendijk and Stockhof work in the possible world tradition and they cast the interpretation of a question as a function which, for every index, returns its true answer. Consider the following example. The semantics given to a question *Does Peter walk?* is a partition of the set of possible worlds into two: those worlds where Peter walks, and those worlds where Peter does not walk. Both *Peter walks* and *Peter does not walk* are possible semantic answers to the question. Each possible world belongs to one or the other of these partitions, so each possible world offers only one *true* answer to the question. This analysis caters for entailment between questions (question *Q* entails question *R* if all true answers to *Q* are also true answers to *R*) and thus explains entailment between coordinated questions. Groenendijk and Stockhof elaborate the basic treatment of yes/no questions; wh-questions are reduced to this basic type. They also provide an interpretation for constituent answers. They assume that modal questions will be analysed at some other pragmatic level.

The work described constitutes the most extensive treatment of the semantics of questions and answers to date. However, in our view, it cannot be directly incorporated in a pragmatic model, for two reasons. First of all, the semantic model assumes completeness of information, and complete mutual awareness of speakers' belief states (but then, why ask questions of one another?). They do at-

tempt to build, from this, an account of how to reason with partial knowledge but, as they work in a traditional extensional framework, this results in clashes with the semantic theory. (In short, what a person knows is a set of possible worlds, namely all those possible worlds that are consistent with his/her beliefs. The semantics of questions is given as a partition over all possible worlds. In an extensional framework - where intensions are derived from extensions - this means that if a person entertains partial beliefs, he/she cannot know the meaning of a question.) Secondly, there may be more than one true answer to a question, and all should be captured by Groenendijk and Stockhof's theory. But how are these answers defined, even computed, from the question? And, as illustrated in the example given in the introduction, even if we know how to generate such answers, how do we define a helpful answer?

3.3 Helpfulness

It is not easy to give a definition of a "helpful" answer off the cuff. Formal semantic theories have little to say on this issue, though some cognitively oriented frameworks have developed useful views.

Relevance Theory [Sperber and Wilson 1986] has given an account of context-sensitivity in communication. It postulates that when people understand, they attempt to maximise relevance - i.e. they pick the context against which the relevance of an utterance is greatest. Relevance, thus, is quantifiable and defined by means of *extent conditions*: an assumption is relevant in a context to the extent that its effects in this context are large and the effort to process it is small.

It should be clear from the onset that the specification of Relevance Theory is not precise enough to be implemented as it stands. There are, however, three principles which are interesting for our purpose. (i) The most relevant context for interpreting a question is that a *Yes*-answer is desired. This helps towards explaining why helpful answers are given at all, and why they occur typically with negative answers. (ii) The selection of relevant contexts is embodied in the human cognitive machinery and ensures that an utterance receives only one interpretation (and not many from which a particular one is selected). Indeed, as shown in the introduction, there may be more than one true answer to a question but only one appropriate one, which must be characterized. (iii) The theory specifies that all contextual effects are explained against the background of assumptions which a person may hold and postulates mechanisms by means of which relevant contexts can be pinned down starting from situational information and the utterance itself.

3.4 Context

The necessity to give a more precise definition of context becomes obvious from the previous sections. Questions can only be answered in context, modals seem to receive different interpretations according to varying contexts, and any cognitively appealing notion of "helpfulness" or "relevance" is stated in terms of contexts. All this ties in with current work in formal semantics which explores tri-partite structures (tying in context with propositional content of utterances) as a basic mechanism for semantic interpretation [Heim 1982, Partee 1988]. However, though current formal semantic theory is steadily increasing the workload of context, its precise nature remains vague. It is not enough to furnish formal semantic interpretations "relative" to some context: a satisfactory approach to a formal but cognitively attractive characterization of "helpful" answers seems to warrant a closer look at the content of conversational backgrounds, their relation to the utterance and its situation, and an appreciation of whether they can be computed.

The insights offered by Relevance Theory may be compatible with formal (and computational) semantic theories, and offer a practical starting point when trying to pin down a fuller notion

of context. In order to investigate this, we need to define our intuitions in an implementable framework. Please note that it is not our intention to attempt a formalisation or an implementation of Relevance Theory, but merely to define an experimental framework capable of handling contexts in order to derive helpful answers to modal and hypothetical questions, albeit exploiting insights from Relevance Theory if possible.

4.0 TOWARDS A THEORY OF PRAGMATICS

It is our intuition that, when people communicate, they know different things, or there would be no point in communicating. Thus it seems that any realistic model of communication must allow for partiality in what agents know. Since agents retain inferential capabilities, we assume their beliefs are consistent. As a consequence, our model represents agents as partial, consistent sets of propositions. The notion of proposition deployed is taken straight from Property Theory [Turner 1987, Chierchia et al 1989, see also Ramsay 1990], a weak first order theory with fine grained intensionality.

Questions are not themselves propositions; they are not associated with truth values. They do, however, entertain a relationship with propositions. In our view, a simple yes/no question embodies a proposition whose truth value is not known to the agent asking the question. An answer to the question is any proposition which, if added to the agent's beliefs, will force truth or falsity of the proposition embodied in the question. This view on answerhood is much looser than the one adopted by Groenendijk and Stockhof in that it allows answers other than true semantic answers. Indeed, any answer will do as long as it allows an agent to *conclude* to the true semantic answer. Thus, a question *Does Peter walk?* may be answered by *Peter sleeps* in this framework (and not just by either *Peter walks* or *Peter does not walk*) as long as that information allows the agent to conclude that *Peter walks* or that *Peter does not walk*. This constrains the agent's reasoning capacity which must now deal with partial information. It also means that agents' beliefs must be subject to revision.

In order to reflect these intuitions in our theory, we extended the language of Property Theory with a predicate which holds of questions, and an operator which, given a proposition, will yield a question. An axiomatisation governs conjunction of questions. A relation of answerhood is defined which holds between a question and its answer (a proposition). The behaviour of this relation is given through axiomatisation of a proof theory.

We adopt a view on modality parallel to Kratzer's: our working hypothesis states that modals are not ambiguous and that any difference in interpretation resides in contextual diversity. We do not, however, try to classify contexts; a hopeless task which is no different to attempting to classify modal ambiguity. *Can* and *must* correspond to the modal operators of *possibility* and *necessity*. Modal constructs are analysed in terms of these operators, a context and a proposition. A context is a collection of propositions, which is a consistent subset of the agent's total beliefs. *Necessity* is true if the negation of the proposition causes a contradiction in the context; *possibility* is true if the proposition can be accommodated within the context without giving rise to a contradiction (i.e. the context can be updated with the proposition).

Questions, whether they are simple or modal, are equally analysed as tri-partite structures comprising an operator, a context and a proposition. For simple yes/no questions the operator is the QuestionTruth predicate (which can be safely stated in Property Theory). For modal questions, the operator is the Question counterpart of the appropriate modal operator. As with Groenendijk and Stockhof, wh-questions are reduced to yes/no questions. It should follow from the above that Groenendijk and Stockhof's results carry over into this model, as the notion of semantic answerhood is preserved (though in an extended framework).

Following Kratzer, conditionals are treated like modals but the context is updated with the antecedent. We are not, however, treating counterfactuals at this stage (i.e. we only treat cases where the context can be updated with the antecedent and where no contradictions occur as a result).

In defining "helpfulness", we take the view of Relevance Theory that a positive answer to the proposition embedded in the question is desirable. As such, *yes*-answers become uninteresting as they are already maximally helpful. *No*-answers, on the other hand, where the proposition cannot be accommodated by the context, can be helpful if they indicate why the proposition is incompatible with a state of affairs, or how the state of affairs might change so that it can be updated with the proposition. In the theory, this information is available from the logic underpinning the answerhood relation relativised to a context. However, this furnishes us with a semantics only. To arrive at some view of how this may interact with pragmatics, the content of contexts must be fleshed out.

Intuition tells us that only one helpful answer is furnished per context. Following Kratzer, and Relevance Theory, we assume that the burden of being helpful and relevant rests with the mechanism which defines the context for an utterance given a situation. Many factors may contribute to this mechanism and it seems reasonable that knowledge of the physical circumstances (i.e. speakers, time, location, etc.) should play a role. The utterance itself must also contribute. As the literature offers no detailed information on how to model the relationship between context and utterance, we have developed an implementation of a context machine which, initially, derived context from lexical information. This implementation was changed and refined in order to attempt to determine experimentally what the requirements for a "context machine" may be.

5.0 THE IMPLEMENTATION

5.1 The Overall Framework

The implementation of the overall framework consists of a parser, a knowledge base, a context machine and a theorem prover. The knowledge base, a consistent collection of propositions, is set up to represent the beliefs of an agent who is to answer questions. For convenience of computation, the items in it are cast as sorted property-theoretic expressions (a sortal hierarchy can be achieved without sorting quantified variables - sorting and closing the world with respect to individuals merely has the effect of rendering the implementation of the first order language decidable). Each knowledge base item is tagged with keys linking the information it contains with words in the lexicon.

The parser, a bi-directional chart parser [Steel and De Roeck 1987] augmented with feature structures, works from an essentially context free rule base where semantic translation rules are paired up with the syntactic statement. The semantic representation delivered by the parser is an expression in Property Theory capturing the structural aspects of question's meaning.

This Property-Theoretic expression is passed to the context-machine. It yields, from the Property-Theoretic expression, a tri-partite structure comprising an operator, a context and a proposition derived from the question. The role of the context machine is to extract from the knowledge base that information which is *relevant* to finding a helpful answer to the question.

The proposition delivered by the context machine is given in the language of the logic K-T [Obeid 1990]. K-T is a propositional, non-monotonic logic which employs Kleene's strong three-valued connectives, and which is extended with two modal operators (the language can be propositional as the knowledge base is sorted and closed). The semantics of the logic are expressed in terms of states of partial information which allow an agent to be uncertain about the truth or falsity of his knowl-

edge, and where possible, to make assumptions on the basis of what is not known to be false. The inference rules of K-T are given in the Appendix.

The propositional content of the input question is set against a suitable subset of the agent's knowledge, i.e. the context. The theorem prover then attempts to prove in the system K-T that the propositional content of the input question follows from the context. This it might achieve monotonically, or non-monotonically with the aid of assumptions. It is the record left in the wake of the proof process in each case, which we interpret in order to provide a helpful answer.

5.2 The Theorem Prover.

The theorem prover is a three valued, modal analogue of a semantic tableau theorem prover [Beth 1962; Jeffrey 1967]. This method performs a case-wise analysis of all models in which the premises (read context) might be true while contradicting the conclusion (read propositional content of the input question). If no such models are found to exist, the theorem is proven. We employ this method because it allows a user absolute access to every stage of the proof process. We then exploit this access in order to find a helpful answer. If a proof succeeds monotonically, the agent's answer is simply *Yes*. If it succeeds by means of one or more assumptions, the answer is of the form *Yes, if ...*, where the body of the *if*-clause is the information that was assumed. Where a proof fails, we have the task of determining the reason why it failed - i.e. which assumptions should be made to yield a *yes*-answer. The proof process constructs a tree of which the branches represent individual models. These models are closely or distantly related to one another according to how much of the proof tree they have in common. A failed proof has one or more models which are consistent, and therefore counterexamples to our intended inference. We are able to compare these consistent models with closely related inconsistent ones. We can then identify the contradiction which is in some sense missing, - i.e. we point to the particular premise or premises which are too weak to support the inference. A helpful answer in this case takes the form *No, unless ...* and the body of the *unless*-clause is composed of the strengthening required in a premise or premises so that the counterexamples would no longer arise. This method remains constant regardless of the actual content of the context. Note that a single answer is always yielded and that the burden of assuring that its content is "helpful" rests entirely on the context machine.

5.3 The Context Machine.

Different implementations of the context selection mechanism have been attempted. Originally, it operated by intersecting that part of the knowledge base which concerns the individuals and relations mentioned in the utterance. In this sense, it relied exclusively on lexical information as the process operated by selecting propositions associated with lexical items reflected as objects in the knowledge base. It used closure on the sortal hierarchy to achieve this. This approach is compatible with Relevance Theory as it can be argued that Encyclopedic Knowledge can be thus implemented. A side effect is lexical disambiguation - different readings of a word are associated with different clusters of information; only compatible information will survive the intersection.

This version was tested on a knowledge base modelling a building site, containing information about buildings, workers, materials and time tables. The domain proved too complex to allow for any conclusions to be drawn: the diversity of objects whose behaviour needed modelling (including some beyond the current state of the art - e.g., mass vs count objects, plurals, time and tense, etc.) was prohibitive. Two other domains were tackled as a consequence: marital relationships and law, and the simple situation of what it takes to drive a car.

Even against simple domains, it became clear that mere reliance on keying lexical information would not be sufficient. The search

space remained large and insufficiently focussed as it included propositions which never contributed to deriving an answer, and a closer interaction between context machine and proof process should be postulated. It seems that the context selection mechanism must have a model of inference. An attempt at such a mechanism was developed.

The context machine Mark II extracts from the knowledge base any information which enable the truth of the proposition associated with the question to be derived. Any implication in the knowledge base with that proposition as a consequent is selected to form part of the context and all rules and assertions which enable the truth of the antecedent of the implication to be derived are also included. Any other rules, which cannot impinge upon the truth of the goal clause, are omitted as they are 'irrelevant' to the proof. In a sense, this selection process anticipates the structure of the proof itself. In the full system, the instantiation of quantified variables in sentences extracted from the knowledge base, is restricted to those individuals mentioned in the question, or relevant to those assertions made about individuals mentioned in the question. This is implemented using the sortal hierarchy. The examples given in Section 5 are derived using this version over a very restricted domain.

Though the results were more satisfactory, the contexts derived in complex domains are still large. Though all information selected plays a part in the overall proof, the search space is uniform for each proof branch. It became clear that a full interaction between the structure of the proof and context selection must be achieved. A third version of the context machine attempted to derive contexts local to particular steps in the proof process. Though incomplete, the experience gained in the attempt convinced us that the selection of 'relevant' contextual information is dynamic. Information pertaining to particular steps in the derivation of an answer should be local to that step and different 'relevant' contexts should be made accessible as the derivation progresses.

6.0 AN EXAMPLE

This section elaborates an example to illustrate (i) the basic theorem prover and (ii) the behaviour of context machines. To simplify the examples, we consider the case where there is only one individual, Anne. The set up concerns finding a helpful answer to *Can Anne drive?* First we present a successful proof, working from an optimal context which yields that Anne can indeed drive. The rules to the theorem prover (K-T [Obeid 1988]) are given in the Appendix. Notice that premises are theorems of the logic and so any premise of form π is logically equivalent to $\sim M \sim \pi$.

[1] M ~ drive(a)	goal	[3] licenced(a)	premise
[2] ownscar(a)	premise	[4] skilltodrive(a)	premise
[5] ownscar(a) & licenced(a) & skilltodrive(a)	-> drive(a)		
			premise
[7] ~M ~ drive(a)	Line 5 - contradiction	[7] [1]	
[6] ~M (ownscar(a) & licenced(a) & skilltodrive(a))			
[8] ~M ownscar(a)	Line 6 - contradiction	[8] [2]	
[9] ~M (licenced(a) & skilltodrive(a))			
[10] ~M licenced(a)	Line 9 - contradiction	[10] [3]	
[11] ~M skilltodrive(a)		0 [4]	

The theorem prover reports that the inference KB \vdash drive(a) is proven by refutation. This we know because each path is inconsistent. The inference was proven monotonically (there was no need for assumptions) and required no sub-proof. The answer here is *Yes*: Anne can drive because she has a licence, she owns a car and she has the required skills.

In the second example, the premise that Anne has a licence is removed. The proof fails to show monotonically that Anne can drive. The system therefore sets out to assume that Anne might have a licence and thus attempts to fill the gap in the agent's in-

formation. Rule R3 (see Appendix) allows us to infer $M \pi$ for any π if we cannot prove $\sim \pi$.

- [1] $M \sim \text{drive}(a)$ goal
- [2] $\text{ownscar}(a)$ premise [4] $\text{skilltodrive}(a)$ premise
- [4] $\text{ownscar}(a) \ \& \ \text{licenced}(a) \ \& \ \text{skilltodrive}(a) \rightarrow \text{drive}(a)$ premise
- [6] $\sim M \sim \text{drive}(a)$ Line 4 - contradiction [6] [1]
- [5] $\sim M (\text{ownscar}(a) \ \& \ \text{licenced}(a) \ \& \ \text{skilltodrive}(a))$
- [7] $\sim M \text{ownscar}(a)$ Line 5 - contradiction [7] [2]
- [8] $\sim M (\text{licenced}(a) \ \& \ \text{skilltodrive}(a))$
- [9] $\sim M \text{skilltodrive}(a)$ Line 8 - contradiction [9] [3]
- [10] $\sim M \text{licenced}(a)$

In this case, if we can assume the premise $M \text{licenced}(a)$ successfully, we can prove the original assertion monotonically. In this context there are no formulae which might affect the truth of $M \text{licenced}(a)$ so our proof succeeds trivially. The answer here is Yes, if Anne has a licence. In the next example, we add explicitly that Anne does not have a licence. We assume that this information is known and does not need a sub-proof.

- [1] $M \sim \text{drive}(a)$ goal [3] $\sim M \text{licenced}(a)$ premise
- [2] $\text{ownscar}(a)$ premise [4] $\text{skilltodrive}(a)$ premise
- [5] $\text{ownscar}(a) \ \& \ \text{licenced}(a) \ \& \ \text{skilltodrive}(a) \rightarrow \text{drive}(a)$ premise
- [7] $\sim M \sim \text{drive}(a)$ Line 5 - contradiction [7] [1]
- [6] $\sim M (\text{ownscar}(a) \ \& \ \text{licenced}(a) \ \& \ \text{skilltodrive}(a))$
- [8] $\sim M \text{ownscar}(a)$ Line 6 - contradiction [8] [2]
- [9] $\sim M (\text{licenced}(a) \ \& \ \text{skilltodrive}(a))$
- [10] $\sim M \text{skilltodrive}(a)$ Line 9 - contradiction [10] [4]
- [11] $\sim M \text{licenced}(a)$

Again, the proof fails monotonically as in the second example. An attempt to hold an assumption that Anne has a licence will, however, fail as it will contradict the premise [3] which states that such an assumption is false. The answer in this case is No, because Anne does not have a licence.

The procedure for dealing with hypotheticals is similar but the context is updated with the antecedent before the proof of the consequent is carried out. Counterfactuals, which would require total revision of the knowledge base, are not treated.

We can use these examples to illustrate the problems faced with selecting the appropriate contexts to yield helpful answers. The earliest version of the context machine would have selected all information associated with domain objects directly related to the words in the sentence (Anne and driving), and all information associated with the sortal hierarchy involving those objects. The union of all these propositions produced a context that was not adequate: only some properties of Anne will affect her driving, and not all knowledge about vehicles will contribute to finding an answer to whether Anne can drive. Intersection of clusters of information obtained by closure on the hierarchy has the side effect of achieving lexical disambiguation, but, in complex domains, it excluded some relevant facts from the context, whilst still including propositions which could never play a role in the proof. A more fine grained approach was needed.

In the second implementation, the context machine selected only those propositions which could lead to a goal. Any implication in the knowledge base with the goal as a consequent is extracted, as are all assertions that contribute to establishing the truth of any of its antecedents (recursively). The proof is established against this context as a whole. Whilst significantly reducing the size of the resulting context as well as focussing its content on what the proof might turn out to be, there are problems with this approach. Imagine the situation where Anne has the skill to drive, she owns a car, but she does not have a licence because she has to pay her fines. She did not pay her fines because she has no money. All this information would be extracted as a total context for answering the question Can Anne drive?

- [1] $M \sim \text{drive}(a)$ goal [3] $\text{skilltodrive}(a)$ premise

- [2] $\text{ownscar}(a)$ premise [4] $\sim M \text{hasmoney}(a)$ premise
- [5] $\text{hasmoney}(a) \rightarrow \text{payfines}(a)$ premise
- [6] $\text{payfines}(a) \rightarrow \text{licenced}(a)$ premise
- [7] $\text{ownscar}(a) \ \& \ \text{licenced}(a) \ \& \ \text{skilltodrive}(a) \rightarrow \text{drive}(a)$ premise
- [8] $\sim M \sim \text{drive}(a)$ Line 7 - contradiction [8] [1]
- [9] $\sim M (\text{ownscar}(a) \ \& \ \text{licenced}(a) \ \& \ \text{skilltodrive}(a))$
- [10] $\sim M \text{ownscar}(a)$ Line 9 - contradiction [10] [2]
- [11] $\sim M (\text{licenced}(a) \ \& \ \text{skilltodrive}(a))$
- [12] $\sim M \text{skilltodrive}(a)$ Line 11 - contradiction [12] [3]
- [13] $\sim M \text{licenced}(a)$
- [14] $\text{licenced}(a)$ Line 6 - Contradiction [14] [13]
- [15] $\sim M \text{payfines}(a)$
- [16] $\text{payfines}(a)$ Line 5 - Contradiction [16] [15]
- [17] $\sim M \text{hasmoney}(a)$

The proof to [13] mimicks that of example 2 above, but now, an attempt to establish whether Anne has a licence requires a sub-proof. The proof fails to close on the assumption that Anne has money. It cannot be inferred non-monotonically that Anne has money (because of [4]). The answer in this case is No, because Anne has no money. Some explanation is due here. Though the answer offered in the last example is "correct", and there might be situations in which it is helpful, it is intuitively arguable that an answer No, because Anne has no licence is more helpful. The point is that this version of the context machine does not cater for the possibility of giving this latter answer under any conditions. From this we conclude that a closer interaction between context machine and proof structure is necessary. A helpful answer should not be confined to the ultimate reasons why the reply is No: the answer should depend upon some measure of "closeness" in contexts. Contrary to the assumptions we made at the start of this project our conclusions lead us to postulate that such a view is indeed necessary to provide a fine grained notion of helpfulness.

We have no treatment of mutual beliefs so far (but Davies [1990] is compatible and promising). We need to extend the logic so it can reason with varying domains if we are to exploit full the intensionality provided by the Property Theory. We have started work on a treatment of time and tense in this framework.

7.0 CONCLUSIONS

We have developed a semantic theory of questions using Property Theory. We have investigated (i) pragmatic answerhood and (ii) modality using an experimental computational framework. We believe that the insights gained from the work have been valuable: they contribute towards our understanding of the requirements for a formally specified and computationally tractable theory of pragmatics which is capable of incorporating insights from cognitively oriented theories. Furthermore, the experiment has pointed out that some of the intuitions underlying Relevance Theory are accurate and useful, especially with respect to context refining strategies necessary for characterising helpful answers.

8.0 REFERENCES

- Ball, R. E.K. Brown, A.N. De Roeck, C.J. Fox, M. Groefsema, N. Obeid and R. Turner (1990) Helpful Answers to Modal and Hypothetical Questions: Final Report, Cognitive Science Centre Memo, University of Essex
- Beth, E.W. (1962) Formal Methods, Dordrecht, Reidel
- Chierchia, G., B. Partee and R. Turner (1989) Properties, Types and Meaning, Dordrecht, Kluwer Academic Publishers
- Coates, J. (1983) The Semantics of the Modal Auxiliaries, London, Croom Helm
- Davies, N. (1990) 'A First Order Logic of Truth, Knowledge and Belief', Proceedings of ECAI 1990

- Groefsema, M., C.J. Fox and N. Obeid (1991) 'Can, May, Must and Should: A division of Labour', Paper accepted at the LAGB, Somerville College, Oxford.
- Groenendijk, J. and M. Stockhof (1984) Studies on the Semantics of Questions and the Pragmatics of Answers, Ph.D. Dissertation, University of Amsterdam
- Heim, I. (1982) The Semantics of Definite and Indefinite Noun Phrases, Ph.D. Dissertation, University of Massachusetts, Amherst (Mass.)
- Jeffrey, R.C. (1967) Formal Logic, its Scope and Limits, London, McGraw-Hill
- Kratzer, A. (1977) 'What 'must' and 'can' must and can mean', Linguistics and Philosophy, Vol 1-3 : 337-355
- Kratzer, A. (1981) 'The Notional Category of Modality', in Eikmeyer and Rieser (eds) Words, Worlds and Contexts, Berlin, Walter de Gruyter
- Obeid, N. (1988) 'A Propositional Logic for Reasoning about Real-Time Situations', in IASTED International Conference, Los Angeles, California.
- Obeid, N. (1990) Partial Models Basis for Non-monotonic Reasoning, Research Note CSM-140, Department of Computer Science, University of Essex
- Palmer, F.R. (1979) Modality and the English Modals, London, Longman
- Palmer, F.R. (1986) Mood and Modality, Cambridge, Cambridge University Press
- Quirk, Q. et al (1985) A Comprehensive Grammar of the English Language, London, Longman
- Ramsay, A. (1990) The Logical Structure of English, London, Pitman Publishing
- Sperber, D. and D. Wilson (1986) Relevance: Communication and Cognition, Oxford, Basil Blackwell
- Steel, S. and A.N. De Roeck (1987) 'Bi-Directional Chart Parsing', in Hallam and Mellish (eds) Advances in AI, London, John Wiley
- Turner, R. (1987) 'A Theory of Properties', Journal of Symbolic Logic, Vol 52, No 2 : 455-472
- Turner, R. (1990), Truth and Modality for Knowledge Representation, Pitman & MIT Press, London.
- Walton, A. (1988), The Pragmatics of English Modal Verbs, Ph.D. Dissertation, University of London.

APPENDIX

Kleene-Turner's (K-T) System [Obeid 1988]

Complete information is hard to obtain, even in the most manageable situations: in most cases, a reasoner does not know everything that is pertinent to the investigation at hand. There are propositions whose truth status cannot be decided. However, most of the classically based non-monotonic formalisms seem to resort to adding intermediary truth values between truth and falsity. This, in fact, is one of the basic and most important features which distinguishes three-valued logics from the classical one. Such a difference is reflected semantically by partial models (partially states of information) for three-valued logics as opposed to possible worlds (complete status of knowledge) for classical logic. In this section, we shall develop the logic K-T.

In K-T a proposition is either accepted as true, accepted as false or not known at all. The basic language L_{K-T} which we shall use is a propositional logic. Starting with primitive propositions T (true), F (false), p, q, r ..., more complicated ones are formed via closure under negation \sim , conjunction $\&$, disjunction \vee , implication \rightarrow and epistemic possibility M. That is, if A and B are well-formed formulae then so are $\sim A$, $A \& B$, $A \vee B$, $A \rightarrow B$ and MA. Let N be the dual of M, i.e. $NA = \sim M$. \sim , \rightarrow , $\&$ and \vee are Kleene's strong connectives. Given A, MA is false if A is false, otherwise MA is true. \sim , $\&$ and M may be taken as primitives. \vee and \rightarrow may be defined in terms of $\&$ and \sim as follows:

Definition 1. $A \vee B = \sim(\sim A \& \sim B)$

Definition 2 $(A \rightarrow B) = \sim A \vee B$

Let $A \leftrightarrow B$ stand for $(A \rightarrow B) \& (B \rightarrow A)$. Note that Kleene's strong implication \rightarrow is not truth functional, i.e. $A \rightarrow B$ is undefined if both A and B are undefined. We also define a truth-functional implication \supset as follows:

Definition 2.3. $(A \supset B) = M(\sim A \& B) \vee \sim A \vee B$. \supset is truth functional in the sense that the truth value of $A \supset B$ is true if both A and B have the same truth value. Let $A = B$ stand for $(A \supset B) \& (B \supset A)$.

Definition 2.4. A model structure for L_{K-T} is $K = \langle B, R, g \rangle$ where B is non-empty set, R is a binary relation on B and g is a truth assignment function g for atomic wffs. The interpretation of R may be thought of as "epistemic possible" extension between states. Given b, b1 are members of B, we shall write b R b1 to mean that the state b1 is an "epistemic possible" extension of the state b.

We employ the notation $K \models g A$ (resp. $K \models \lrcorner g A$) to mean that A is accepted as true (resp. false) in K with respect to g. For convenience, reference to g will be omitted except when a confusion may arise. Let A, B be wffs; then the truth \models and the falsity $\models \lrcorner$ notions are recursively defined as follows:

Definition 2.5.

- (i) $K, b \models T$
- (ii) $K, b \models p$ iff $g(b, p) = \text{true}$ for p atomic
- (iii) $K, b \models A \& B$ iff $K, b \models A$ and $K, b \models B$
- (iv) $K, b \models \sim A$ iff $K, b \not\models A$
- (v) $K, b \models MA$ iff $(\exists b1 \in B) (b R b1 \text{ and } K, b1 \not\models \sim A)$
- (i') $K, b \models \lrcorner F$
- (ii') $K, b \models \lrcorner p$ iff $g(b, p) = \text{false}$ for p atomic
- (iii') $K, b \models \lrcorner A \& B$ iff $K, b \models \lrcorner A$ or $K, b \models \lrcorner B$
- (iv') $K, b \models \lrcorner \sim A$ iff $K, b \models A$
- (v') $K, b \models \lrcorner MA$ iff $(\forall b1 \in B) (b R b1 \text{ then } K, b1 \models \sim A)$

The logic K-T is the smallest set of L_{K-T} which is closed under the following axiom schemas and inference rules. We shall write $\vdash_{K-T} A$ to mean that A is a "theorem" of K-T.

Axiom Schemas:

- (a1) $A \supset (B \supset A \& B)$
- (a2) $A \supset (B \rightarrow A)$
- (a3) $A \& B \supset A \mid A \& B \supset B$
- (a4) $A \rightarrow B \supset [(B \rightarrow C) \supset (A \rightarrow C)]$
- (a5) $\sim \sim A \equiv A$
- (a6) $\sim(A \& B) \equiv (\sim A \vee \sim B)$
- (a7) $A \rightarrow MA$

Inference Rules:

Modus ponens (Mp) for \supset together with (R1), (R2) and (R3).

- (R1) From $\sim A \vee B$ infer $\sim MA \vee B$
- (R2) From $A \supset B$ infer $MA \rightarrow MB$
- (R3) -----

$\vdash MA$