

A Process-Oriented, Intensional Model of Knowledge and Belief

by

Robert F. Hadley

School of Computing Science
Simon Fraser University

1. INTRODUCTION

Within the cognitively related disciplines (including AI) there are three dominant approaches towards the modeling of epistemic states. For convenience we call these the logistic, syntactic, and intensional approaches. The *logistic* approach seeks to discover a *logic* of belief (knowledge) which would enable one to derive, *a-priori*, the totality of things an agent believes (or knows), given the formulas explicitly written in the agent's belief base.¹ The *syntactic* approach (as Levesque [1984] dubs it) identifies an agent's beliefs with just those formulas explicitly written in the agent's belief base -- the agent has no other beliefs. By contrast, the *intensional* approach identifies beliefs with abstractions, e.g., with *thought contents* or propositions. On this account the same belief may be expressed by means of a limited variety of distinct syntactic objects. As will emerge in section 1.3, our present concerns lie with the intensional model, and with how intensions are related to the *processes* of inference, belief acquisition, and belief retraction. However, we now briefly describe the logistic and syntactic approaches, and consider how these differ from the present work.

1.1 The Logistic Approach

Those who adopt the logistic approach are concerned with the discovery of *inference rules*, such that an agent's explicit beliefs are *closed under* the application of these rules. Typically, the epistemic logics proposed do not require an agent's beliefs to be closed under *all* valid inference rules (full logical omniscience is not required), but closure under *certain* inference rules is still required. Thus, these logics typically require that agents each possess an infinity of beliefs, or at least believe an infinity of distinct formulas. (cf. [Levesque, 1984; Halpern & Moses, 1985; Lakemeyer, 1987]). Some have questioned the motivation for introducing logics which represent agents as believing an infinity of formulas (e.g., [Fagin & Halpern, 1985]), but the underlying rationale may be that such logics employ a non-standard sense of belief, which yet has application in the context of finite agents. For example, the sense of belief appropriate to [Levesque, 1984] *might* be such that "*p* is believed" means approximately that one could *eventually come to believe p* by a tractable computational process. However, whether or not one thinks that epistemic logics (of the "infinite belief set" variety) require us to adopt a non-standard sense of belief, such logics usually are not concerned with certain important *process-related* aspects of belief. For example, with limited exceptions² such logics do not address the fact that an agent acquires and *changes* beliefs over time, or that these processes often involve inference, testing beliefs for consistency, retraction, etc. Doubtless, most epistemic logics were never intended to address those aspects of belief, but for certain applications one desires a *formal model* of how these cognitive processes relate to the belief states of the agent. An agent *X* who knows what another agent *Y* believed an hour ago may wish to form a reasonable conjecture about what *Y* believes *now*. If so, *X* needs to know what *processes* must be assumed to exist in order to arrive *even* at a reasonable conjecture about *Y*'s present beliefs. Of course, some aspects of belief change are addressed by truth maintenance systems (e.g., [de Kleer, 1986; Doyle, 1979]), but TMSs are concerned with *procedures* for manipulating belief sets, rather than with formal models which interrelate the concepts of belief, knowledge, inference, retraction, etc. By contrast, we are here concerned with formalizations of epistemic states, and the processes which affect those states.

1.2 The Syntactic Approach

¹Hereafter, "belief" shall be taken to include knowledge, unless the context requires that the distinction be explicit.

²Konolige, [1985] presents an epistemic logic which does address the process of belief *introspection*. However, Konolige's system is not intended to address most other process-related aspects of belief.

Now, apart from epistemic *logics* which posit an infinity of beliefs, certain *syntactically-oriented* belief models do address the fact that beliefs can be acquired and retracted by *processes* (cf. [Haas, 1985]). In general, those who adopt the syntactic approach identify beliefs with particular syntactic objects, so that, for example, an agent might believe the formula $p \vee q$ but not $q \vee p$, or believe "Tadpoles swim" but not "Pollywogs swim". Now, many reject the syntactic approach entirely, on the grounds that it is too fine-grained. The syntactic approach arguably fails to account for the fact that *some* syntactically distinct formulas seem to express precisely the same thought, and therefore ought to have the same belief status (cf. [Levesque, 1984] ; [Hadley, 1986]).

Fagin and Halpern [1985] present what might be viewed as a *liberalized* syntactic belief model. The Fagin-Halpern (F-H) model identifies an agent's beliefs with a subset of formulas of which the agent is aware. This particular subset (which satisfies special conditions) *may* be taken to include formulas the agent will only *come to be* aware of at some future time. In this respect the F-H model is less restrictive than those which restrict an agent's beliefs to formulas presently written in the agent's belief base. However, apart from the awareness requirement, the F-H model is not concerned with the empirical preconditions of belief (e.g., that a sophisticated robot might derive a formula from existing beliefs, but not "commit to" the formula until a "consistency check" with other beliefs is complete).

There is no reason to suppose that the F-H model was ever intended to deal with the empirical preconditions of belief, but an agent who wishes to reason (or conjecture) about what a given robot *currently* believes will require a model of belief which makes explicit many of these preconditions. Such a model will need to recognize the fact that agents revise and retract beliefs, and that they sometimes derive formulas which they *never* accept, even though the derived formula is indeed a valid consequence of the agent's existing beliefs.

1.3 The Intensional Approach

The intensional approach is founded upon the premise, widespread among philosophers of language, that belief and knowledge must be understood as a relation between an agent and an abstract *intension* (proposition or thought). Recent work in epistemic modeling, which adopts an intensional stance, includes that of Wilks and Ballim [1987], and Maida [1986]. Wilks and Ballim are concerned primarily with the *generation* of nested beliefs (especially atypical beliefs) and Maida concentrates upon epistemic problems which arise between agents whose reasoning strategies are unknown to each other. Nevertheless, these authors share our assumption that agents believe a particular symbolic representation (sentence, formula) insofar as they believe the intension *expressed* by that representation.

The thesis that a declarative sentence expresses an abstract sense or intension, whose structure is a *function* both of the sentence's syntax and of the meaning of its parts, is a view which dates back at least to Frege. Modern philosophical descendents of Frege, including Montague [1970] and Lewis [1976], have considerably refined and modified Frege's view, but Frege's underlying idea that the sense (intension) of a sentence is determined by a compositional process, involving both syntax and meaning functions, persists. The precise definition of 'intension' varies from theory to theory, but a sentence's intension is usually taken to be a *function* which, given a possible world as argument, returns that truth value of the sentence in that possible world.

One notorious difficulty with intensional semantics is that when intensions are merely construed as functions in the set-theoretic sense, then every tautology, ϕ , turns out to have the same intension (or meaning), since each valid ϕ must be assigned the value 'true' in each possible world. (Recall that ϕ 's intension is a function from possible worlds to truth values). Thus, when intensions are regarded *merely* as functions, intensions are too coarse-grained to serve as meanings or "thought contents". In an attempt to circumvent this difficulty, Lewis [1976] defines sentence meaning as a function both of a sentence's *parse tree* and its intension. In doing so, he quickly disposes of the problem concerning tautologies. However, many believe Lewis's modification results in a theory of meaning which is excessively fine-grained, since under his modification virtually *every* syntactically distinct pair of sentences (e.g., $p \vee q$ vs. $q \vee p$) is assigned a different meaning. In [Hadley, 1988] a simple modification of Lewis's view is proposed to remove this difficulty. The modification creates a degree of granularity for intensions which allows us to say that *structurally* non-isomorphic formulas (including tautologies) differ in intension, but formulas which are both logically equivalent and structurally isomorphic in a graph-theoretic sense, have the same intension or meaning. For example, on this account $p \vee q$ and $q \vee p$ have the same intension since they are logically equivalent and their parse trees are isomorphic.

Many other *compositional* theories of sentence meaning exist, of course, and each has its own degree of granularity. Cresswell [1985] proposes a theory of sentence meaning with an attractive degree of granularity, as does Gawron [1986]. For present purposes, we shall remain neutral among these differing theories. We shall simply identify intensions with sentence meanings, and the reader may select the theory of meaning which suits his/her purpose. What is relevant for our purposes is that each of the aforementioned theories proposes a *canonical* system for representing meanings or intensions, such that *all sentences which are synonymous* (according to the semantic theory) *are mapped onto a single canonical representation*. Moreover, each of the semantic theories includes a *principled* set of rules for translating (or compiling) an expression into its canonical meaning representation. Questions about synonymy may be answered by comparing the canonical representations into which the respective expressions are "compiled" by the translation rules. Questions whether two distinct expressions can represent the same belief will in turn be decided by ascertaining whether the distinct expressions are synonymous.³

1.4 The Proposed Model

In what follows an axiomatic system is presented which aims to formalize many of the interrelationships among the epistemic concepts, and related notions such as retracting, questioning, or committing to a belief. The model is designed to capture human-like epistemic concepts, and is intended to be of use to those who seek *formal* models of epistemic *processes*. Unlike many who adopt the *logistic* approach, our overriding concern is not with finding a tractable model, but with providing a conceptual map which is sensitive to the empirical contingencies of belief and knowledge.⁴ Moreover, because our focus is primarily upon *cognitive* fidelity, it is hoped that the axiomatic structure will make vivid just how many assumptions one must make to conjecture that an agent has some *actual* belief at a given time.

The epistemic model to be described is axiomatized in first-order logic. In spite of this, the model permits the representation of *nested* beliefs, and permits agents to reason about the beliefs of themselves and others. Self-reference is possible within the language of the model; thus, the potential for paradoxes of self-reference arises. A method (due to Kripke [1975]) for eliminating the contradictions which self-reference can engender is described in section 3.

In section 2 the axioms of the formal model are presented. The axioms modify and generalize the *intensional* model of belief proposed in [Hadley, 1988] (which included neither *knowledge*, the representation of nested epistemic propositions, nor self-referential belief statements). As will become apparent, the model proposed here (hereafter 'NIM' for "nested intensional model") requires an agent to know or believe *only* those formulas which are intensionally identical with formulas explicitly present in the agent's belief base. Thus, NIM does not entail logical omniscience. Moreover, NIM does not require us to suppose that agents believe every formula they have validly *derived* from other beliefs.

2. AN INTENSIONAL EPISTEMIC MODEL

The axiom set described below (section 2.1) permits agents to reason about their own epistemic states and those of others. Thus, the language of our model includes nested (embedded) epistemic formulas. However, the language should also permit us to say that different agents believe different formulas at different times. For this reason, epistemic assertions should involve at least three separate parameters. The representation of assertions involving multiple parameters is straightforward in first-order logic (FOL), and for this reason (among others) FOL is chosen. However, the representation of *nested* epistemic propositions in FOL is less straightforward.

Perlis (1985) describes a method, involving quotation, unquotation, and concatenation of expressions, which permits the representation of nested propositions in an extension of FOL. However, this method can be unwieldy for levels above the first level of embedding. For this reason, we here adopt the simple expedient of treating the epistemic operators (K and B) as skolem functions. For example, we will suppose that the function $K(i, f, t)$ simply returns the *proposition* that agent i knows formula f at time t . When we wish to assert this proposition, we simply attach the predicate 'true', as in $true(K(i, f, t))$. (We assume lower case for predicates and variables, and upper case for

³The reader is referred to [Hadley, 1988] for a detailed discussion of the compilation process and the individuation of intensions.

⁴There is no reason to suppose that all relevant aspects of belief and knowledge can be expressed in tractable formal systems. We know, for example, that humans engage in certain forms of reasoning which are intractable. For more on this, see Perlis [1985].

functors). We do not attempt to define the *truth* predicate within our axiom set, but assume it to be interpreted from without (cf. section 3). Moreover, we need not *always* explicitly display the truth predicate; rather we stipulate that any functor written in bold face is preceded by an implicit truth predicate. So, to say that (for all agents, formulas, and times) if it is true that a formula is known to be known, then that formula is true, we may write:

$$\mathbf{K}(i, \mathbf{K}(j, f, t_1), t_2) \rightarrow \text{true}(f).$$

(Unless otherwise indicated, assume variables are universally quantified with widest possible scoping). Note that the foregoing assumptions approximate the Fregian view that when a formula occurs within an epistemic context, it denotes its *sense*, and otherwise, it denotes its truth value.

2.1 Axioms of the Model

Before presenting the axioms of NIM, it may prove helpful to make explicit a few more background assumptions. In the following we assume that every agent possesses a belief base (BB), which contains a finite collection of explicit formulas which the agent regards as true. Any axiom which the agent accepts must be represented in the BB by at least one formula. Synonymous (intensionally equivalent but syntactically distinct) formulations of such axioms need not be present. Beliefs which are *caused* in the agent via processes not normally considered inferential (e.g., via sense perception) will be regarded, for present purposes, as axioms acquired at various times. In addition, the BB may contain formulas which the agent has explicitly inferred and "accepted". We assume that all formulas in a BB belong to an artificial language, but we make no assumptions about the particular language and inference rules which the agent X employs. We do *not* assume that X's axioms form a consistent set, or that X never makes a fallacious inference.

We further assume that agents *may* employ strategies for maintaining consistency, and may, as a consequence, retract previously held beliefs. The retraction process may, or may not be ideally rational. Finally, we do not assume that an agent enters a sentence into his BB as soon as it is inferred. There is a sense of 'infer' according to which one believes anything one has inferred, but here we use 'infer' to mean merely 'derive according to accepted rules (which may include heuristics)'. Agents may treat newly inferred formulas as "candidates" for belief, and enter an examination or "questioning" phase during which the candidate is tested for consistency with all or some of the agent's other beliefs.

We now present the axioms of NIM. As will be apparent, NIM contains an unusual number of predicates whose interpretation is underconstrained by the formal structure. This is due, we believe, to the fact that it deals with *aspects* of belief which are simply ignored in most *axiomatic* treatments. Our contention is that one simply cannot form reasonable conjectures about an agent's belief states over time, unless these additional aspects are explicitly considered. For clarity, the intended interpretations of each axiom and predicate are discussed as the axioms are introduced.

The expression $\mathbf{B}(i, f, t)$, which occurs in most of the axioms, may be interpreted as *it is true that i believes f at time t*. The predicate $\text{in-bb}(i, f, t)$ may be rendered *formula f is explicitly present in i's BB at time t*.

1. $\text{in-bb}(i, f, t) \rightarrow \mathbf{B}(i, f, t)$

Axiom 1 tells us that all formulas explicitly present in *i*'s BB at time *t* are among *i*'s beliefs at *t*. In axioms 2 and 3, below, the intended interpretation of $\text{sense}(f, s)$ is: *s is the intension of formula f*. The variable *s* ranges over the class of intensions. Thus, axiom 2 states that an agent who believes a formula *f* believes all formulas synonymous with *f*. Axiom 3 makes an analogous statement about knowledge. Recall that, for artificial languages at least, the synonymy of formulas may be ascertained by compiling the formulas into their canonical intensional representations, and checking the results for identity.

2. $[\text{sense}(f_1, s) \wedge \text{sense}(f_2, s)] \rightarrow [\mathbf{B}(i, f_1, t) \equiv \mathbf{B}(i, f_2, t)]$

3. $[\text{sense}(f_1, s) \wedge \text{sense}(f_2, s)] \rightarrow [\mathbf{K}(i, f_1, t) \equiv \mathbf{K}(i, f_2, t)]$

The following axiom expresses the widely accepted view that whoever knows *f* must at least believe *f*, and *f* must be true.

$$4. \quad \mathbf{K}(i, f, t) \rightarrow \mathbf{B}(i, f, t) \wedge \text{true}(f)$$

Note that, if we assume $\text{true}(K(\text{Mary}, K(\text{John}, p, t1), t2))$ (which we abbreviate as $\mathbf{K}(\text{Mary}, K(\text{John}, p, t1), t2)$), then we may use axiom 4, together with conjunctive simplification, to derive $\text{true}(K(\text{John}, p, t1))$. This in turn enables us to derive $\text{true}(p)$, by analogous reasoning. Axiom 5 simply states that an agent who retracts a formula ceases to believe the formula. This, together with axiom 4, entails that the formula would not be known. We may use axiom 2 to show that the agent would also cease to believe all synonymous formulas. Axioms 1 and 5 together entail that a retracted formula ceases to be in the belief base.

$$5. \quad \text{retracts}(i, f, t) \rightarrow \neg \mathbf{B}(i, f, t)$$

In the following axiom, $\text{questioning}(i, f, t)$ may be interpreted as *i begins at time t to investigate whether formula f is consistent with other beliefs*. The underlying idea is that, if *i* has begun to question whether *f* should be retained, then *i* is not presently committed to *f*'s truth. We assume that those automated reasoning systems, which employ belief revision strategies, enter a phase in which the consistency of a set of formulas is investigated before any retractions are made. A formula which is being questioned may be one which has been singled out for special investigation. However, we do not insist that 'retracts' and 'questioning' be interpreted just as suggested here, and no attempt is made to formalize these notions. If a given system lacks a questioning mechanism, or even a retraction mechanism, then the 'questioning' predicate (or 'retracts') may be assigned the value 'false', which would render the presence of these predicates innocuous.

$$6. \quad \text{questioning}(i, f, t) \rightarrow \neg \mathbf{B}(i, f, t)$$

The intended interpretation of axiom 7 is that any formula which has just been inferred by an agent, from a subset of that agent's beliefs, is to be regarded as a candidate for belief by that agent. In the following, $\text{just-inferred}(i, f, s, t)$ may be read *i inferred f from set s at time t*. $\text{Mem}(w, s)$ may be read as *w is a member of set s*, and $\text{cand}(f, i, t)$ is rendered as *f is a candidate for i at t*.

$$7. \quad [\text{just-inferred}(i, f, s, t) \wedge \forall w \{ \text{mem}(w, s) \rightarrow \text{in-bb}(i, w, t) \}] \rightarrow \text{cand}(f, i, t)$$

Axiom 8 tells us that, under normal circumstances, beliefs persist. More specifically, the axiom tells us that if *i* believes a formula at t_1 , then *i* also believes the formula at all later times, unless *i* retracts or questions the formula at some intervening time. Since the interpretation of 'retracts' is left to the user's discretion, 'retracts' may be taken to include some form of spontaneous forgetting or other memory loss. Nothing here forces us to restrict retraction to deliberate actions. In 8, $\text{later}(t_1, t_3)$ may be read as t_3 is later than t_1 , and $\text{inc-between}(t_1, t_2, t_3)$ may be read t_2 lies inclusively between t_1 and t_3 .

$$8. \quad [\mathbf{B}(i, f, t_1) \wedge \text{later}(t_1, t_3) \wedge \neg \exists t_2 \{ \text{inc-between}(t_1, t_2, t_3) \wedge [\text{retracts}(i, f, t_2) \vee \text{questioning}(i, f, t_2)] \}] \rightarrow \mathbf{B}(i, f, t_3)$$

The preceding eight axioms are intended to be *descriptive* of our concept of belief. They relate the concept of belief to cognate concepts in a manner intended to be useful to agents who reason about the beliefs of an artificial agent. The following axiom, by contrast, is *prescriptive*. It is intended to convey the conditions under which a candidate formula (in the sense introduced by axiom 7) *should* become an active belief. The term ' $\text{successful}(i, f, t_2)$ ' should be taken to mean that at time t_2 formula *f* has just been successfully scrutinized by agent *i*'s questioning mechanism, and has been found acceptable. The term ' $\text{should-be-inserted}(i, f, t_2)$ ' may be read as, formula *f* should be added to *i*'s BB at time t_2 . Thus axiom 9 states that a formula which is a candidate for *i* at t_1 , and which has successfully passed the questioning phase at a later time t_2 should be added to *i*'s BB at t_2 . Axiom 1 would then ensure that *f* is a belief at t_2 .

$$9. \quad [\text{cand}(f, i, t_1) \wedge \text{later}(t_1, t_2) \wedge \text{successful}(i, f, t_2)] \rightarrow \text{should-be-inserted}(i, f, t_2)$$

Now, since the interpretation of 'questioning' may vary according to the intended application of the model, the interpretation of 'successful' should vary accordingly. If a particular agent altogether lacks a questioning mechanism, we may simply assign 'true' to 'successful', since there would be no tests which a candidate formula should pass, beyond being a candidate. Of course, questions may arise about how exhaustive a questioning or consistency testing process must be before we may assign 'success' to a formula. However, these are practical issues which must be decided in particular applications.

2.2 Formal Properties of the Model

We now consider some of the formal properties of NIM. First, note that the elements of NIM, with minor syntactic alterations, may be represented in any standard first-order predicate calculus, known to be sound and complete. We use 'PC' to denote any such logic.

Proposition 1. NIM is consistent.

The consistency of NIM may be checked by transforming each of its axioms into clause form, using a standard clause-form transformation algorithm. The NIM axioms are consistent just in case the resulting clauses are jointly consistent. One may easily verify that these clauses are simultaneously satisfiable in a universe of one individual. If the universe contains but one individual, all variables in the clauses may be replaced by a unique constant denoting this individual. It is then trivial to assign truth values to the resulting grounded literals so that each of the clauses is rendered true by the given interpretation.

Definition. Let **T** be a first-order theory formed by adding to PC an encoding of the axioms of NIM in PC.

Proposition 2. **T** is sound and complete.

Proof. The soundness of **T** follows from the fact that its rules of inference are exactly those of PC, and all axioms in **T** are consistent. If PC contains axioms, they are universally valid, and thus consistent with the consistent set NIM, which was added to PC to obtain **T**. The completeness of **T** follows immediately from the completeness of PC. The inference rules of PC ensure that all logical consequence of the NIM axioms encoded in PC are derivable.

3. DISCUSSION

We now return to a difficulty previously mentioned, namely, that NIM includes an explicit truth predicate. Since our model does not include a *definition* of the truth predicate, it is not vulnerable to the kind of internal inconsistencies which Tarski [1936] describes. Nevertheless, paradoxes involving self-reference can still arise on a *given interpretation* of the predicates in NIM. It may be objected, therefore, that the existence of potentially paradoxical sentences renders NIM defective. However, Kripke has clearly demonstrated [1975] that *any* language rich enough to permit the *ascription* of truth to sentences is vulnerable to paradox. He further demonstrates that this vulnerability in no way presupposes syntactic or semantic ill-formedness, but can arise for perfectly innocent looking sentences when unusual empirical conditions exist. Of course, one can always banish truth ascriptions in a special-purpose language, but if we are ever to have an adequate model of *human cognition*, we will need to represent propositions such as, "What Mary believes is true".

Since we are ultimately concerned with adequate cognitive models, we have included a truth predicate. Thus, paradoxical sentences may arise. However, both Kripke [1975] and Perlis [1985] have shown that the inclusion of such sentences need not render a first-order language inconsistent. Kripke proposes a complex theory of truth according to which paradoxical sentences are simply not assigned truth values. The underlying intuition of Kripke's theory is that any ascription of truth to a sentence involves a *presupposition* that the sentence is grounded. (Roughly, a grounded sentence is one whose truth evaluation "bottoms out", either immediately or ultimately, in the evaluation of sentences which do not ascribe truth or falsity to other sentences.) On the rare occasions when the presupposition of groundedness is unsatisfied no truth value is assigned. Kripke goes on to argue (following Kleene [1952]) that the existence of sentences without truth values does not require that we abandon any of the usual laws of logic. For example, since " $p \vee \neg p$ " still holds for all cases when *p* has a truth value, we may retain the law. When *p* lacks a truth value, the disjunction is *not false*; rather it also lacks a truth value.

4. CONCLUSION

A generalized first-order model of belief and knowledge has been presented which simultaneously avoids logical omniscience and the excessively fine-grained "syntactic approach". The model, NIM, is intensionally based, and sanctions inferences involving nested epistemic attitudes, with different agents and different times. NIM makes explicit many (if not all) of the empirical assumptions an agent must make before concluding that another agent believes (knows) something *at a given time*. Because the model is axiomatized in FOL, it may be used by any standard, first-order theorem prover. Moreover, NIM provides agents with a conceptual map, interrelating the concepts of knowledge and belief and a number of cognate concepts, such as 'infers', 'retracts', and 'questions'. Because the model builds upon the concept of an *intension*, whose degree of granularity is left open, the range of formulas which are permitted to *express the same belief* is left to the user of NIM. For example, *one could* define "same intension" so that all logically equivalent formulas would have the same intension, and thus (by axiom 2) would express the same belief. However, if one wishes to retain an intuitive, human-like concept of belief, one

HADLEY

should individuate intensions so that (roughly) only those formulas which we intuitively regard as *synonymous* are permitted to have the same intension, and thus express the same belief. We have suggested several semantic theories which yield (approximately) the required criterion of synonymy.

REFERENCES

- Cresswell, M.J., (1985) *Structured Meanings: The Semantics of Propositional Attitudes*, MIT Press, Cambridge, MA.
- de Kleer, J. (1986) "An Assumption-based TMS", *Artificial Intelligence*, Vol. 28, pp. 127-162.
- Doyle, J. (1979) "A Truth Maintenance System", *Artificial Intelligence*, Vol. 12, pp. 231-272.
- Fagin, R., and Halpern, J., (1985) "Belief, Awareness, and Limited Reasoning", Proceedings of IJCAI, Los Angeles, pp. 491-501.
- Frege, G. (1952) "On Sense and Reference", in *Translations from the Philosophical Writings of Gottlob Frege*, Oxford.
- Gawron, J.M. (1986) "Types, Contents, and Semantic Objects", *Linguistics and Philosophy*, Vol. 9, No. 4, pp. 427-476.
- Haas, A., (1985) "Possible Events, Actual Events, and Robots". *Computational Intelligence*, Vol. 1, No. 2, pp. 59-70.
- Hadley, R.F. (1988) "Logical Omniscience, Semantics and Models of Belief", *Computational Intelligence*, Vol. 4, No. 1.
- Hadley, R.F., (1988b) "A Default-Oriented Theory of Procedural Semantics", *Cognitive Science*, to appear.
- Hadley, R.F., (1986) "Fagin and Halpern on Logical Omniscience: a Critique with an Alternative", Proceedings Sixth Canadian Conference on Artificial Intelligence, Montreal.
- Hobbs, J.R., and Rosenschein, S.J., (1978) "Making Computational Sense of Montague's Intensional Logic", *Artificial Intelligence*, Vol. 9, pp. 287-306.
- Kleene, S.C., (1952) *Introduction to Meta-Mathematics*, Wolters-Noordhoff Publishing and North-Holland Publishing, New York.
- Konolige, K., (1985) "A Computational Theory of Belief Introspection", Proceedings of IJCAI, Los Angeles, pp. 503-508.
- Lakemeyer, G., (1987) "Tractable Meta-Reasoning in Propositional Logics of Belief", Proceedings of IJCAI, Milan, pp. 401-408.
- Levesque, H.J., (1984) "A Logic of Implicit and Explicit Belief", FLAIR Tech. Report #32, Fairchild, Palo Alto, Calif.
- Lewis, D. (1976) "General Semantics", in *Montague Grammar*, (ed.) Partee, B., Academic Press, New York.
- Maida, A.S. (1986) "Introspection and Reasoning about the Beliefs of Other Agents", *Proceedings of the Cognitive Science Society*, pp. 187-195.
- Montague, R., (1970) "Universal Grammar", *Theoria* 36.
- Perlis, D., (1985) "Languages with Self-Reference I: Foundations", *Artificial Intelligence*, Vol. 25, pp. 301-322.
- Tarski, A. (1936) "Der Wahrheitsbegriff in den formalisierten Sprachen", *Studia Philos.*, Vol. 1, pp. 261-405.
- Wilks, Y. & Ballim, A. (1987) "Multiple Agents and the Heuristic Ascriptions of Belief", Proceedings of IJCAI, Milan, pp. 118-124.