The detection and representation of ambiguities of intension and description

Brenda Fawcett and Graeme Hirst
Department of Computer Science
University of Toronto
Toronto, Ontario
CANADA M5S 1A4

Abstract

Ambiguities related to intension and their consequent inference failures are a diverse group, both syntactically and semantically. One particular kind of ambiguity that has received little attention so far is whether it is the speaker or the third party to whom a description in an opaque third-party attitude report should be attributed. The different readings lead to different inferences in a system modeling the beliefs of external agents.

We propose that a unified approach to the representation of the alternative readings of intension-related ambiguities can be based on the notion of a *descriptor* that is evaluated with respect to intensionality, the beliefs of agents, and a time of application. We describe such a representation, built on a standard modal logic, and show how it may be used in conjunction with a knowledge base of background assumptions to license restricted substitution of equals in opaque contexts.

1. Introduction

Certain problems of ambiguity and inference failure in *opaque contexts* are well known, opaque contexts being those in which an expression can denote its *intension* or underlying concept rather than any particular *extension* or instance. For example, (1) admits two readings:

- (1) Nadia is advertising for a penguin with whom she could have a long-term meaningful relationship.

On the *transparent* (or *extensional* or *de re*) reading, there is some particular penguin that Nadia is after:

- (2) Nadia is advertising for a penguin with whom she could have a long-term meaningful relationship, whom she met at a singles bar last week and fell madly in love with, but lost the phone number of.

On the *opaque* (or *intensional* or *de dicto*) reading, Nadia wants any entity that meets her criteria:

- (3) Nadia is advertising for any penguin with whom she could have a long-term meaningful relationship.

On this reading, the rule of *existential generalization* fails; that is, we cannot infer from (3), as we could from (2), that:

- (4) There exists a penguin with whom Nadia could have a long-term meaningful relationship.

Another rule of inference that fails in opaque contexts is *substitution of equals*; (5) and (6) do *not* permit the conclusion (7):

- (5) Nadia believes that the number of penguins campaigning for Greenpeace is twenty-two.
- (6) The number of penguins campaigning for Greenpeace is forty-eight.
- (7) $\nRightarrow$ Therefore, Nadia believes that forty-eight is twenty-two.

Although these facts are familiar, little research has been done on how a practical NLU system can detect and resolve intensional ambiguities (which can occur in many constructions besides the 'standard' examples; see Fodor 1980, Fawcett 1985), and control its inference accordingly. The same is true of certain other complications of opaque contexts that are of special relevance to systems that use explicit representations of knowledge and belief. In particular, the interaction between intensional ambiguities

and the beliefs of agents has not been studied. The present work is a first step towards rectifying this.

2. Attributing descriptions

Previous linguistic systems that dealt with opaque contexts, such as that of Montague (1973), have taken a God's-eye view, in the sense that the speaker and listener are assumed to have perfect knowledge, as are, in certain ways, the people of whom they speak. No account is taken of the limits of the knowledge or beliefs of the agents involved.

To see that beliefs are a complicating factor, consider the following sentence, usually considered to be two ways ambiguous — transparent or opaque:

- (8) Nadia wants a dog like Ross's.

These ambiguities, however, cross with an ambiguity as to which agent the description *a dog like Ross's* is to be attributed: to the speaker, or to Nadia (the agent of the verb of the sentence). This gives a total of four possible readings. To see the four cases, consider the following situations, all of which can be summarized by (8):

- (9) *Transparent reading, agent's description:*
Nadia sees a dog in the pet store window. "I'd like that dog," she says, "It's just like Ross's." The speaker of (8), who need not be familiar with Ross's dog, reports this.
- (10) *Transparent reading, speaker's description:*
Nadia sees an animal in the pet store window. "I'd like that," she says. Nadia is not aware of it, but the animal is a dog just like the one Ross owns. The speaker of (8), however, knows Ross's dog (and believes that the listener also does).
- (11) *Opaque reading, agent's description:*
Nadia feels that her life will be incomplete until she obtains a dog. "And the dog that would be perfect for me," she says, "Is one just like the one that Ross has." The speaker of (8), not necessarily familiar with Ross's dog, reports this.
- (12) *Opaque reading, speaker's description:*
Nadia feels that her life will be incomplete until she obtains a pet. "And the pet that would be perfect for me," she says, "Is a big white shaggy dog, with hair over its eyes." Nadia is not aware of it, but Ross owns a dog just like the one she desires. The speaker of (8), however, knows Ross's dog (and believes that the listener also does).

The agent's-description readings permit the inference that Nadia *believes* that she (either intensionally or extensionally) wants a dog like Ross's; the other readings do not. Making the distinction is thus crucial for any system that reasons about the beliefs of other agents, such systems being an area of much current concern in artificial intelligence (*e.g.*, Levesque 1983, Fagin and Halpern 1985).

Another complicating factor is the time at which a description is to be applied. The above readings assumed that this was the time of the utterance. The intensional readings, however could be referring to the dog that Ross will get or (not included in the examples below) once had:

- (13) *Opaque reading, agent's description, future application:*
Nadia has heard that Ross will buy a dog. Wanting one herself, and trusting Ross's taste in canines, she resolves to buy whatever kind he buys.
- (14) *Opaque reading, speaker's description, future application:*
Nadia finds English sheepdogs attractive, but none are available. She therefore intends to purchase some other suitably sized dog and spend her weekend gluing long shaggy hair onto it. Nadia is not aware of it, but Ross owns a dog just like the one she wants to end up with. The speaker, knowing Ross's dog, can describe Nadia's desire as that of having an object that will at some future time be describable as a dog like Ross's.

The description in an intensional reading may also be used to refer to different entities at different times.

- (15) *Opaque reading, agent's description, repeated application:*
Ross buys a new type of dog every year or so. Desperately wanting to keep up with canine fashion, Nadia declares her intent to copy him. Whatever dog Ross has at any given time, Nadia wants to have the same kind.

We have not been able to find an example in which repeated application of the speaker's description gives a natural reading. Extensional readings always seem to refer to the present time.² Thus, there are at

²It may be objected that an extensional future-application reading is also possible. This would be like (14), except that Nadia has some particular dog in mind for the cosmetic alterations. If we allow Nadia to use this method repeatedly upon a particular dog, then an extensional reading corresponding to (15) would be

least seven readings for *Nadia wants a dog like Ross's*.³

3. Other intensional ambiguities and inference failures

There are other kinds of intension-related inference failures besides those mentioned in the previous sections. For example, some opaque contexts forbid inferences from postmodifier deletion, while others permit it. Both readings of (16) entail the less specific (17) (which preserves the ambiguity of (16)):

- (16) Nadia is advertising for a penguin that she hasn't already met.
- (17) Nadia is advertising for a penguin.

However, the same cannot be done with (18):

- (18) Nadia would hate for there to be a penguin that she hasn't already met.
- (19) $\nRightarrow$ Nadia would hate for there to be a penguin.⁴

The examples above have all involved explicit or implicit propositional attitudes and such contexts are apparently necessary for ambiguities of attribution of description and the associated possible inference failure and for problems of postmodifier deletion. However, there are many other kinds of context in which other intension-related ambiguities and inference failures can occur. For example, existential generalization can also fail in contexts of similarity and possibility:

- (20) Nadia is dressed like a creature from outer space.
- (21) $\nRightarrow$ There is a creature from outer space whom

derived. That is, Nadia wants her particular dog to once or repeatedly become like Ross's dog. However, we don't see these readings as distinct from (14) and (15); Nadia's desire is clearly towards the goal of having a dog that matches a particular description, rather than towards that of owning a particular dog.

³Hofstadter, Clossman, and Meredith (1982) analyze a similar sentence for the case where the speaker and the agent are the same, *I want the fastest car in the world*, and derive five readings where we predict four. However, their two extensional readings are identical in our analysis, as they differ only in how many separate descriptions the agent has for the entity.

⁴This example is based on one of Fodor's (1980: 188). Fodor claims that postmodifier deletion is never valid in an opaque context; as example (17) shows, this claim is too strong. The problem in (19) seems to be that *would hate* means *wants not*, and the deletion is invalid in the scope of a negation.

Nadia is dressed like.

- (22) It is possible that a creature from outer space could interrupt your lecture at the most inconvenient moment.
- (23) $\nRightarrow$ There is a creature from outer space who could possibly interrupt your lecture at the most inconvenient moment.

The kind of semantic irregularities that we are discussing are thus found in a large and syntactically diverse set of linguistic constructs. (See Fawcett (1985) for a long list, with discussion and examples.) Many seem to display idiosyncratic semantic features that could necessitate a broad range of operators in a representation, destroying any apparent homogeneity of the class. It is our suggestion, however, that these constructs can be processed in a uniform way. We argue that the diversity among the constructs can be accounted for by evaluating *descriptors* according to intensionality, agents, time, and states of affairs. Introducing the concept of a descriptor preserves the homogeneity of the class, while the dimensions along which descriptors may vary provide enough detail to differentiate among the particular semantics of the constructs.

4. The descriptor representation

In this section we introduce a representation designed to capture the different possible readings of opaque constructions. In developing the representation, we have tried to move away from previous approaches to intensionality, such as that of Montague (1973), which use truth conditions and meaning postulates, and which take no account of the beliefs or knowledge of agents. Influenced by recent work on situation semantics (Barwise and Perry 1983, Lespérance 1986) and belief logics, we have aimed for a more 'common-sense' approach.

In the representation, we take an *intension* to be a finite representation of those properties that characterize membership in a class, and by a *descriptor* we mean a non-empty subset of the elements of an intension (in practice, often identical to the complete intension). A descriptor provides access either to the intension of which it is a part or to its extension. This eliminates the need of explicitly listing all the known properties of an entity; only properties

relevant to the discourse situation are mentioned.

The representation is described in detail in Fawcett (1985); below we give a short description of the main points, and some examples of its use.

The representation is based on conventional temporal modal logic. The general form of a completed sentential clause is a proposition of the form

$(term\text{-}list) \langle predication \rangle$.

The term-list, which can be empty, contains all the quantified terms except those which are opaque with respect to agents or time; the predication expresses the main relation among the various entities referred to. The intention is that the term-list provides the information to identify referents in the knowledge base, and the main predication asserts new information to be added to it. Usually the argument positions of the predication will be filled by bound variables or constants, introduced previously in the term-list. However, within temporal operator or agent scopes, argument positions may instead contain quantified terms. Term-list-predicate pairs may be nested inside of one another.

Quantified terms arise from noun phrases. They have the general form

$(Det\ X: R(X))$

where *Det* is a quantifier corresponding to the explicit or implicit determiner of the noun phrase, *X* is the variable introduced, and *R(X)* indicates restrictions on *X*. In the examples below, we restrict ourselves to only three quantifiers — *indef*, *def*, and *label*, introduced by indefinite descriptions, definite descriptions, and proper nouns respectively.⁵

To this formalism, we add the following:

- The agent scope marker $\wedge$.
This marker can apply to a formula or term to indicate that any embedded descriptors must be evaluated with respect to the beliefs of the agents involved (that is, mentioned so far) at the point where the scope of $\wedge$ begins. The speaker is assumed to always be available as an agent, and descriptors outside the scope of $\wedge$ are attributed only to the speaker.

⁵For simplicity, we treat names as extensional in our examples. However, there is nothing to prevent an opaque treatment, in which the different agents are thinking of different individuals with the same name.

- The intensional abstractor **int-abs**.
The formula
 $\text{int-abs}(C, (Quant\ Var : Description))$
asserts that the quantified term *Var* is to have an intensional referent (*i.e.*, an individual or universal concept), which is returned in *C*. If *C* is subsequently used, then its referent is a universal (generic) concept, which we do not discuss in this paper; see Fawcett (1985) for details. If *Var* is used instead, then the referent is an individual concept. (Without **int-abs**, use of *Var* refers to an extension.)
- Descriptors.
The notation [*d X*] indicates that the properties *d* are being used as a descriptor of entity *X*. Thus its intensionality, time of application, and agent must be considered. (Variables over such descriptors are permitted, so we can manipulate them independently of the entities to which they might refer.)

Thus, opacity with respect to agents and opacity with respect to time are both treated as scope ambiguities, while intensionality is marked as a binary distinction. In general, all quantified terms are left-extrapolated to the outermost term list. Those quantified terms marked as intensionally ambiguous may be prefixed by **int-abs**. Those quantified terms originating within the scope of the agent scope marker $\wedge$ may remain inside its scope and be evaluated relative to the agents available at that point. Similarly, those quantified terms originating in the scope of the temporal operators **F** and **P** (future and past) may stay inside their scope, thus indicating a future or past application of the descriptor.

The following example shows the representations of the first four readings of (8) (*i.e.*, those with the description applied at the time of the utterance), and an extensional counterpart. (In the examples, the quantifier *indef* corresponds to the English determiner *a*, and the quantifier *label* is used for proper nouns. The structure of the descriptor *dog-like-Ross's*, orthogonal to our concerns here, is not shown.)

- (24) *Transparent reading, agent's description:*
There is a dog Nadia wants, and she describes it as being like Ross's dog.
(label *Y* : Nadia)
 $\langle \text{want } Y, \wedge (\text{indef } X : [\text{dog-like-ross's } X]) \rangle$

- (25) *Transparent reading, speaker's description:*
 There is a dog Nadia wants, and the speaker describes it as being like Ross's dog.
 (label Y : Nadia)
 (indef X : [dog-like-ross's X])
 <want Y , $\wedge X$ >
- (26) *Opaque reading, agent's description:*
 Nadia wants something she describes as being a dog like Ross's.
 (label Y : Nadia)
 <want Y ,
 $\wedge$ int-abs (C , (indef X : [dog-like-ross's X]))>
- (27) *Opaque reading, speaker's description:*
 Nadia wants something that the speaker describes as being a dog like Ross's.
 (label Y : Nadia)
 int-abs (C , (indef X : [dog-like-ross's X]))
 <wants Y , $\wedge X$ >

Note that the fourth reading has no representation in a conventional first-order modal language. For comparison, here is a non-opaque sentence of the same structure.

- (28) Nadia buys a dog like Ross's.
 (label Y : Nadia)
 (indef X : [dog-like-ross's X])
 <buy Y , X >

Within the scopes of the opaque operators **F**, **P**, and $\wedge$, special checks must be made before standard inference rules can apply.⁶ We do not assume that *all* arguments are intensional; we favour a policy towards intensional scopes of "introduce when required" to minimize the amount of extra processing needed. Our use of the symbol $\wedge$ is quite different from that of Montague. For Montague, $\wedge x$ denotes an object that is intensional. We instead use this notation to delimit the agent scope of an opaque construct; descriptors in x are potentially ascribed to any of the agents preceding the $\wedge$ marker.

Our approach to determiners is a compromise between other common approaches. The first, common in computational linguistics, is to represent determiners by three-place quantifiers of the general

⁶This is analogous to the restricted rules that Montague presents for substitution of identicals and lambda conversion in his intensional logic (Dowty, Wall, and Peters 1981: 165). We seek a more flexible scheme that, rather than prohibiting inference, restricts its use to certain special cases.

form

$$det(x, R(x), P(x))$$

where x is the variable introduced, R is the restriction on the variable, and P is the new predication on the variable. This reflects observations of Moore (1981) and others that determiners rarely have a direct correlation with the existential and universal quantifiers of first-order logic. In many of the meaning representations used with logic grammars (Dahl (1981), for example), determiners provide the basic structure of the meaning representation formula. The determiners are translated into quantifiers and are all left-extraposited (to be later scoped relative to one another on the basis of some reasonably simple set of rules). As a result, the main predication of a clause will always be nested in the rightmost predication position.

Another approach focuses more on the main verbs by first translating them into predicates, and subsequently finding appropriate fillers for their arguments that contain the necessary quantifiers. However, this does not allow a convenient way to represent relative scoping ambiguities. Montague combines the two approaches. All quantifiers introduce two predicates: a restriction predicate and a main predication as in

$$\lambda R \lambda P (\exists x (R\{x\} \text{ AND } P\{x\})),$$

which translates the determiner *a*.

Our approach is a compromise. Quantified terms consist of a variable and restriction, but do not incorporate the main predication. All quantified terms (except those that are opaque with respect to time or agent) are left-extraposited and assimilated into a single list structure followed by a single main predication.

5. Substitution of equals

Given our descriptor logic, we can now turn to the question of when substitution-of-equals inferences can and can't be made.

The failure of substitution of equivalent phrases appears to be a *gradable* notion; the degree of substitution allowed varies with the type of construct under consideration. We can think of a scale of substitutivity, with the lower bound being a strictly de

dicto reading in which no substitutions are permitted and the upper bound a strictly de re reading in which co-extensional phrases can be substituted in any context.

For example, sentences that refer directly to the form of the expression admit no substitution:

- (29) The Big Bopper was so called because of his size and occupation.
- (30) The Big Bopper was J. P. Richardson.
- (31) $\nrightarrow$ J. P. Richardson was so called because of his size and occupation.

In sentences of propositional attitude, certain descriptors can be substituted for, provided the content of the proposition, relative to the speaker and the hearer, is not affected. It is easy to recognize such cases, but not always easy to specify what exact criteria determine terms that are interchangeable. Consider:

- (32) Nadia thinks that the Queen of England is a lovely lady.
- (33) Nadia thinks that Queen Elizabeth is a lovely lady.
- (34) Nadia thinks that the titular head of the Church of England is a lovely lady.

The assumption is that since the filler of the role *Queen of England* is not likely to change within the time of the conversation and the speaker, the hearer, and Nadia are all aware of who fills that role, it is acceptable to substitute the filler for the role and vice versa. Thus, sentence (33) can be inferred from (32). But to substitute the phrase *the titular head of the Church of England*, as in (34), seems to attribute more knowledge to Nadia than was in the original statement.

The problem of substitution in opaque contexts stems from the failure to recognize how descriptors relate, and not, as in classical logical approaches, from the failure of expressions to be “co-intensional”. The emphasis should be on identifying the relation between descriptors with respect to appropriate agents rather than on co-intensionality alone; in most cases co-intensionality is too strong a condition for substitution. Rather, the *background assumptions* of the discourse determine whether a substitution of one descriptor for another is

permitted.

A typical substitution replaces the target descriptor, $d1$, with an equivalent descriptor, $d2$, from the background assumptions, but otherwise preserves the form of the target sentence, i.e.,

$$\text{RESULT} = \text{TARGET} [d1/d2].^7$$

To see whether a descriptor substitution is valid in an opaque context, three factors must be checked in the following order: the *intensionality* of the descriptor, the *time of reference* of the descriptor, and the *agents* of the descriptor. We must establish the “level” of each factor in the target sentence and then determine whether the background assumptions authorize substitutions at that level. That is, we must relate the intensionality, time, and agent of the *descriptor equivalence* asserted in the background assumptions to those of the target descriptor, and then assert the intensionality, time, and agent of the descriptors in the resulting clause (after any substitutions).

The background assumptions will have already been derived from earlier input (in a manner described by Fawcett 1985, section 5.5) and assimilated into the system’s general knowledge base. In order to compare descriptors in the target to descriptors in the background assumptions, we extract the relevant aspects from the representation of each, and express them explicitly by the use of the following *descriptor predicates*, which can then be used to query the knowledge base.

- **desc** ($a, e, d1$).
Ascribes a particular descriptor to an individual; “agent a would use the descriptor $d1$ to describe the entity e ”.
- **label** ($a, c, name$).
Indicates that the label $name$ is known by agent a to be a label for the (individual) constant c .
- **time** ($t, e, d1$).
Asserts that descriptor $d1$ describes entity e at time t .

As an example, consider the four readings of this sentence in which the description is applied at the time of utterance:

⁷Not all substitutions are of this form; see Fawcett 1985, section 5.4.

(35) Nadia wants the fastest car in the world.

(i) *Extensional reading, speaker's description:*
(label Y: Nadia)
(def X: [fcw X]) <want Y, ^X>

(ii) *Intensional reading, speaker's description:*
(label Y: Nadia)
int-abs (C, (def X: [fcw X])) <want Y, ^X>

(iii) *Extensional reading, agent's description:*
(label Y: Nadia)
<want Y, ^ (def X: [fcw X])>

(iv) *Intensional reading, agent's description:*
(label Y: Nadia)
<want Y, ^**int-abs** (C, (def X: [fcw X]))>

(fcw stands for the descriptor *fastest-car-in-the-world*.)

Table I lists some different possible background assumptions. We will show the different effects of each. Background assumption I asserts the co-extensionality of the descriptors *fastest car in the world* and *Ross's Jaguar 300*, while assumption II asserts co-intensionality of the descriptors. Assumptions III and IV express the same equivalences, and, additionally, knowledge of them is also attributed to Nadia.

When the beliefs of agents (other than the listener) are not involved, the following rule licenses certain substitutions of equivalents:

- If the target descriptor is intensional⁸ then co-intensional or definitionally equivalent descriptors in the background assumptions may be substituted.

Background assumptions I and II thus allow substitutions in readings (i) and (ii), as shown in table II. (For simplicity, the quantifier

(label Y: Nadia)

is omitted from each example.)

When attribution of descriptions is involved, as in readings (iii) and (iv) of (35), we must determine whether the other agents are (believed by the listener to be) aware of the equivalence. The general rule for substituting descriptors which are ambiguous with respect to descriptive content is this:

- If the assertion of descriptor equivalence in the background assumptions in the listener's knowledge base is part of the knowledge base of the agent to

⁸In this rule, the descriptor must not be generic. Rules for generics (universal concepts) are described in Fawcett 1985, section 5.4.

TABLE I
BACKGROUND ASSUMPTIONS

I	The fastest car in the world is Ross's Jaguar 300.
II	The fastest car in the world (always) is a Jaguar 300.
III	Nadia believes that the fastest car in the world is Ross's Jaguar 300.
IV	Nadia believes that the fastest car in the world is a Jaguar 300.

TABLE II
SUBSTITUTIONAL INFERENCE

(i) + I	Nadia wants Ross's Jaguar 300. (def X: [ross's-jag300 X]) <wants Y, ^X>
(i) + II	Nadia wants a Jaguar 300. (def X: [jag300 X]) <wants Y, ^X>
(ii) + I	<i>No substitution possible.</i>
(ii) + II	Nadia wants a Jaguar 300. int-abs (C, (def X: [jag300 X])) <wants Y, ^X>
(iii) + III	Nadia wants Ross's Jaguar 300. <wants Y, ^ (def X: [ross's-jag300 X])>
(iii) + IV	Nadia wants a Jaguar 300. <wants Y, ^ (indef X: [jag300 X])>
(iv) + III	<i>No substitution possible.</i>
(iv) + IV	Nadia wants some Jaguar 300. <wants Y, ^ int-abs (C, (indef X: [jag300 X]))>

whom the target descriptor is ascribed, then the descriptor can be substituted in the target. The resulting clause will have the substituted descriptor attributed to the same agents as the descriptor in the original target.

Reading (iii) requires a co-extensional descriptor that Nadia is aware of. Background assumptions III and IV both provide such a descriptor. Reading (iv) also requires a descriptor that Nadia is aware of, but it must be co-intensional with the target descriptor; only assumption IV provides such a descriptor which can then be substituted. The results are shown in table II.

Substitution rules for other intensional constructs, and details of interactions between rules, can be found in Fawcett (1985, section 5.4).

6. Implementation

We have implemented a prototype system that incorporates the ideas discussed above. The system is written in Prolog, and is built on top of Popowich's SAUMER formalism for syntactic and semantic rules (Popowich 1984, 1985).

7. Plans and goals

Now that we have looked at the problem of detecting these ambiguities and representing the possible readings, the next step is to study how the ambiguities may be resolved, and what factors influence the preference for one reading over another. We expect that in most cases pragmatic factors will be central, although there may be default preferences in some constructions. In addition, another member of our group, Diane Horton, is studying the interaction between agents' descriptions and the presuppositions of a sentence (Horton 1986).

Acknowledgements

This paper is based on thesis work by the first author under the supervision of the second, who also wrote the paper. The authors acknowledge helpful discussions with each other, Diane Horton, and Hector Levesque, and financial support from IBM, the Natural Sciences and Engineering Research Council of Canada, and the University of Toronto. They are also grateful to Nick Cercone and Fred Popowich for making the SAUMER system available to them.

References

- BARWISE, Jon and PERRY, John (1983). *Situations and attitudes*. Cambridge, MA: The MIT Press / Bradford Books, 1983.
- DAHL, Veronica (1981). "Translating Spanish into logic through logic." *American journal of computational linguistics*, 7(3), 149-164.
- DOWTY, David R; WALL, Robert E; and PETERS, Stanley (1981). *Introduction to Montague semantics* (Synthese language library 11). Dordrecht: D. Reidel, 1981.
- FAGIN, Ronald and HALPERN, Joseph Y (1985). "Belief,

awareness, and limited reasoning: Preliminary report." *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, Los Angeles, August 1985. 491-501.

FAWCETT, Brenda (1985). *The representation of ambiguity in opaque constructs*. MSc thesis, published as technical report CSRI-178, Department of Computer Science University of Toronto, October 1985.

FODOR, Janet Dean (1980). *Semantics: Theories of meaning in generative grammar* (The language and thought series). Cambridge, Mass.: Harvard University Press, 1980.

HOFSTADTER, Douglas R; CLOSSMAN, Gary A; and MEREDITH, Marsha J (1982). "'Shakespeare's plays weren't written by him, but by someone else of the same name.' An essay on intensionality and frame-based knowledge representation." Bloomington, Indiana: Indiana University Linguistics Club, November 1982.

HORTON, Diane (1986). *Incorporating agents' beliefs in a model of presupposition*, MSc thesis, Department of Computer Science, University of Toronto, forthcoming (June 1986).

LESPEANCE, Yves (1986). "Toward a computational interpretation of situation semantics." *Computational intelligence*, 2(1), February 1986.

LEVESQUE, Hector (1983). "A logic of implicit and explicit belief." *Proceedings of the National Conference on Artificial Intelligence (AAAI-83)*, Washington, D.C., August 1983, 198-202.

MONTAGUE, Richard (1973). "The proper treatment of quantification in ordinary English." [1] In: Hintikka, Kaarlo Jaakko Juhani; Moravcsik, Julius Matthew Emil and Suppes, Patrick Colonel (editors). *Approaches to natural language: Proceedings of the 1970 Stanford Workshop on Grammar and Semantics*. Dordrecht: D. Reidel, 1973. 221-242. [2] In: Thomason, Richard Hunt (editor). *Formal Philosophy: Selected Papers of Richard Montague*. New Haven: Yale University Press, 1974. 247-270.

MOORE, Robert C (1981). "Problems in logical form." *Proceedings of the 19th Annual Meeting, Association for Computational Linguistics*, Stanford, June 1981, 117-124.

POPOWICH, Fred (1984). "SAUMER: Sentence analysis using metarules." Technical report 84-2, Laboratory for Computer and Communications Research, Simon Fraser University, Burnaby, B.C., Canada. August 1984.

POPOWICH, Fred (1985). "The SAUMER user's manual." Technical report 85-4, Laboratory for Computer and Communications Research, Simon Fraser University, Burnaby, B.C., Canada. March 1985.