

**THE REPRESENTATION OF AMBIGUITY
IN OPAQUE CONSTRUCTS**

Brenda Fawcett

**Technical Report CSRI-178
April 1986**

**Computer Systems Research Institute
University of Toronto
Toronto, Canada
M5S 1A1**

The Computer Systems Research Institute (CSRI) is an interdisciplinary group formed to conduct research and development relevant to computer systems and their application. It is jointly administered by the Department of Electrical Engineering and the Department of Computer Science of the University of Toronto, and is supported in part by the Natural Sciences and Engineering Research Council of Canada.

THE REPRESENTATION OF AMBIGUITY
IN OPAQUE CONSTRUCTS

by

Brenda Fawcett

Department of Computer Science
University of Toronto
Toronto, Ontario, Canada
October 1985

A Thesis submitted
in conformity with the requirements
for the degree of Master of Science

Copyright © 1985 Brenda L. Fawcett

Acknowledgements

My primary debt of gratitude is to my supervisor, Graeme Hirst. His diligence and efficiency at commenting and criticizing ideas put before him was most helpful. Especially valuable for me was his exceptional patience and persistent enthusiasm for my work. A sincere thank you to him.

I am also grateful to two families for their unrelenting support. The Master and Fellows of Massey College were invaluable sources of encouragement and friendship, and the variety of activities at the College provided memorable alternatives to thesis-writing.

My own family played a more passive but practical role. The idea of their daughter still being a student after six or so years, I'm sure is somewhat puzzling for them. I thank them for their co-operation and understanding.

I wish to thank my fellow graduate students, in particular, Allan DeNooy and Diane Horton, for enduring with me in my more foolhardy moments. Diane also helped me enormously in revising and typesetting the thesis as a technical report. I am further indebted to those who accepted my all-too-frequent squash challenges and thus sustained my rather fervent addiction.

Finally, I thank Hector Levesque for his time and effort as a second reader, and acknowledge IBM and the Natural Sciences and Engineering Research Council for their financial support which made student life, at worst, rather comfortable.

For my grandmothers.

Abstract

A knowledge of intensions, which are used to designate concepts of objects, is important for natural language processing systems. Certain linguistic phrases can refer either to the concept of an entity or to the entity itself. To properly understand a phrase and to prevent invalid inferences from being drawn, the system must determine the type of reference being asserted. We identify a set of "opaque" constructs and suggest that a common mechanism be developed to handle them. Our approach is intensional in that it maintains a distinction between extensional and intensional reference, yet it does not adhere rigidly to the strict, logical definition of intension.

To account for the ambiguities of opaque contexts, noun phrases are translated into *descriptors*. It must be made explicit to whom the descriptor is ascribed and whether its referent is non-specific or specific. Similarly, sentential constituents should be treated as *propositions* and evaluated relative to conjectured *states of affairs*. As a testbed for these ideas we define a Montague-style meaning representation and implement the syntactic and semantic components of a moderate-size NLP system in a logic programming environment.

One must also consider how to disambiguate and interpret such a representation with respect to a knowledge base. Much contextual and world knowledge is required. We attempt to characterize what facilities are necessary for an accurate semantic interpretation, considering what is and is not available in current knowledge representation systems.

Table of Contents

Chapter 1 Introduction	1
1.1 The Concept of Intentions	1
1.2 Intentions for Natural Language Processing	3
1.3 Purpose and Scope	4
1.4 Thesis Overview	6
Chapter 2 Linguistic Motivations for Intentions	7
2.1 Syntactic Classification	7
2.2 Semantic Considerations for Opaque Contexts	10
2.2.1 Inference Failures	10
2.2.1.1 Existential Generalization	10
2.2.1.2 Substitution of Equivalents	12
2.2.1.3 Other Logical Inferences	13
2.2.2 Presupposition	13
2.2.3 Deletion of Modifying Phrases	14
2.3 Noun Phrases, the prime target for ambiguity	14
2.4 Factors Influencing Ambiguity	15
2.4.1 Definiteness of Noun Phrases	16
2.4.2 Degree of Description of Noun Phrases	16
2.4.3 Verb Tense and Aspect	16
2.4.4 Verb Type	16
2.4.5 Surrounding Linguistic Context	16
Chapter 3 Semantic Classification	18
3.1 The Use and Evaluation of Descriptors	18
3.1.1 Referents for NP Descriptors	20
3.1.2 Equivalence Among NP Descriptors	22
3.1.3 Labels, Special-Purpose Descriptors	23
3.2 States of Affairs	24
3.3 Semantic Classification	25
3.3.1 Agentless Sentential Opaque Contexts	25
3.3.1.1 Objective Attitudinals	26
3.3.1.2 Temporal Phrases	27
3.3.1.3 Deontic Phrases	28
3.3.2 Sentential Opaque Contexts With Agents	28
3.3.2.1 Propositional Attitudes	28
3.3.2.2 Reported Behaviour	29
3.3.3 Intensional Transitive Verbs	29

3.3.4 Direct Quotations and Titles	30
3.3.5 Intensional Modifiers	31
3.4 Concluding Remarks	31
Chapter 4 Descriptor Logic	33
4.1 Montague Semantics Revisited	33
4.2 Descriptor Logic as a Meaning Representation	33
4.2.1 Basic Entities	33
4.2.2 General Form for NL Constituents	34
4.2.2.1 Correspondence with Linguistic Constituents	36
4.2.3 Candidate Set Of Operators	36
4.2.3.1 Propositional Operators	36
4.2.3.2 Term Predicates	37
4.2.3.3 Other Things	37
4.3 Examples	38
4.4 Discussion of Decisions	38
4.4.1 Opaque Scopes	38
4.4.2 Determiner Translation	38
4.4.3 Adjectives and Adverbs	39
4.4.4 Iterated Opaque Contexts	39
4.4.5 Non-Opaque Scoping	40
4.5 Descriptor Relations	41
4.5.1 Merging of Descriptors	42
Chapter 5 Substitution of Equals	43
5.1 Range of Substitution	43
5.2 The Substitution Phenomenon	44
5.2.1 A Detailed Example	44
5.3 Establishing the Background Assumptions	45
5.3.1 Copular Relations and their Descriptions	46
5.4 Minimal Requirements for Substitution	50
5.4.1 Intensionality of the Descriptor	50
5.4.2 Time of Descriptor Reference	53
5.4.3 Agents of Descriptors	54
5.5 Extracting Descriptor Information from the MR	57
5.5.1 Knowing Labelled Referents and Substitution	58
5.5.2 Querying the Knowledge Base	58
Chapter 6 NLP Systems and Opacity	60
6.1 Compatability of Logic Grammars and Montague Semantics	60
6.2 Specifications for an Implementation	61
6.3 Refinements of the MR for the NLP System	62
6.4 Types of Objects in the Implemented MR	62
6.5 The Knowledge Base Interpretation	64
6.5.1 Adequacy of Available KR Schemes	64
6.5.2 Intensionality	65

6.5.2.1 Asserting Relations Among Such Nodes	68
6.5.2.1.1 Ordinary Classification	70
6.5.2.1.2 Generic Classification	70
6.5.2.1.3 Extensional Equivalence	71
6.5.2.1.4 Intensional Equivalence	72
6.5.2.1.5 Definitional Equivalence	73
6.5.2.1.6 Extensional Predication	75
6.5.2.1.7 Intensional Predication	76
6.5.3 Agents in the Knowledge Base	77
6.5.4 Time and States of Affairs	78
6.5.4.1 Choosing the Current State of Affairs	79
6.6 Conclusions and Suggestions for Further Research	79
Bibliography	81
Appendix A Classification of Noun Phrases	86
1 Definite Descriptions	86
2 Indefinite Descriptions	86

CHAPTER 1

Introduction

The relative ease with which humans use natural language to communicate complex ideas has intrigued scholars from a wide variety of disciplines. The "Natural Language Problem"¹ most recently has been addressed by computer scientists in the field of computational linguistics. *Pure* linguists seek to describe the regularities of language and the rules governing them. In contrast, *computational* linguists strive to actually apply such principles both for the understanding² and generation of natural language using a computer. Thus linguistic theories are combined with methods of computer science, notably those of Artificial Intelligence (AI), to construct efficient computational models for natural language (NL).

Natural language processing (NLP) systems apply rules which define the syntax, semantics, and pragmatics of a given NL. The *syntactic* component translates the surface structure of a sentence into a possible constituent tree structure; the *semantic* component transforms these constituent structures into some useful meaning representation; the *pragmatic* component specifies a mapping from the meaning representation to a deeper, unambiguous representation which takes account of non-linguistic factors such as world knowledge and principles of discourse. Our particular concern is the design of the syntactic and semantic components of a NLP system able to handle the class of sentences containing *opaque constructs*. These constructs involve verbs such as *believe*, *think*, *imagine* in which the substitution of phrases designating the same entity does not preserve the truth value of the construct. Our approach will incorporate the notion of intensions into the meaning representation. Schubert and Pelletier (1982) have argued that an intensional representation is not appropriate for NLP. They favour a more "conventional" logic without intensional operators.

The choice between intensional and conventional translations is as yet unclear. Given that the conventional approach appears to have certain advantages, it is worth finding out where it leads; but we are not irrevocably committed to this approach. (Schubert and Pelletier 1982:29)

Thus, further motivation is to provide counter-arguments to those of Schubert and Pelletier and so present the case *for* an intensional treatment.

1.1. The Concept of Intensions

intension : the suggesting of a meaning by a word apart from the thing it explicitly names or describes.

(Webster's Third New International Dictionary : 481)

The term *intension* intuitively can be paraphrased as "the concept of" or, more common in AI, "the description of". Given a linguistic expression, say "the Prime Minister of Canada", we can identify two potential referents of the phrase. On one interpretation it

¹ The exact definition of the "Natural Language Problem" varies among disciplines (*e.g.*, philosophers, linguists, and psychologists). The pervasive question seems to be "What makes natural language useful for communication and what are the mechanisms underlying it?"

² "Understanding" is used here in a very narrow sense, namely to mean that some representation is constructed to encode the meaning of a phrase.

refers to a particular entity, at the present time the individual named Brian Mulroney. Alternatively, this same phrase can be used to refer to the *concept* of the entity, that is, the abstract bit of mentalese which represents the entity whose properties include "being leader of the elected government in Canada", "living at 24 Sussex Drive", and the like. The first reading refers to the *extension* of "the Prime Minister", while the second refers to what will be considered the *intension*. In certain contexts both readings are possible and it is important to decide which one is intended.

Intensionality has been considered in two related yet distinct fields. The first is within linguistics and philosophy as part of their theories of language and thought. The questions raised in this context concern the epistemological status of intensions, their role in human reasoning, and how we (as humans) use them to interpret the world about us.³ In general, AI researchers will not be concerned with these issues. In particular, this thesis is concerned with what communication takes place, but not with any cognitively real account of how it occurs.

The second field in which intensions have been extensively studied is logic, especially in model-theoretic semantics as typified by the works of Frege (1870), Carnap (1947), and Kripke (1963). These logicians were interested in formalizing the notion of intension as a well-defined object in semantic theory. Their motivations stemmed from the inadequacy of extensional logics (*e.g.*, first-order logics) to account for the failure of standard inference rules in certain linguistic contexts. Consider the following:

(1-1) Susan *believes that* the Morning Star is the Morning Star.

(1-2) The Morning Star is the Evening Star.

(1-3) Susan *believes that* the Morning Star is the Evening Star.

The truth of sentences (1-1) and (1-2) does not allow us to conclude that (1-3) is true, even though constituents with the same referent have been intersubstituted, namely "the Evening Star" and "the Morning Star" both referring to the planet Venus. This often-quoted, classic example of Frege's exhibits the failure of substitution of equals in *belief*-contexts. Frege also noted the apparent failure of the principle of compositionality which states that the semantic value of a sentence should be a function of the semantic values of its parts. It predicts that the truth value of the whole sentence should be dependent on the truth value of any embedded propositions⁴. It is clear that such is not the case. In example (1-3), whether the embedded clause "the Morning Star is the Evening Star" expresses a proposition that is actually true or not has no bearing on the truth value of the *belief*-sentence containing it.

Such examples led Frege to distinguish between the "sense" and the "reference" of a linguistic expression. The sense (or *Sinn*) of an expression concerns the mode of presentation; that is, the way in which an object, if any, is referred to by a given phrase. The reference (or *Bedeutung*) then refers to the actual object designated by the phrase.

The ideas of Frege were made more rigorous by Carnap (1947) and Kripke (1963) in their modal and possible worlds semantic systems. "Sense" or intension was expressed as a

³ The linguists' and philosophers' definition of intension (synonymous here with "concept") relies heavily on intuition. The most common usage is to identify a "concept" with "that sort of mental representation which expresses a property and is expressed by an open sentence. For example, the concept DOG can be thought of as

- a) an image of some generic dog, or
- b) expressing the property of dogness, or
- c) the mental representation involved in the utterance of the open sentence *is a dog* " (Fodor 1981 : 260).

⁴ Here a truth conditional analysis for NL is assumed whereby each sentence is associated with a proposition and the referent of that proposition is a truth value.

function from possible world indices into entities or sets of entities. Later Church (1951) gave an axiomatic treatment of such entities, and Montague (1974) developed an intensional logic system with a model-theoretic interpretation specifically applied to natural language. While some connection between these formal objects of logic and the philosophical notion of intension is assumed to exist, the correspondence is intuitive and not at all necessitated by the logical theories themselves.

We consider our approach *intensional* for two reasons:

- (1) The underlying motivation is to provide a definition of expressions not just for a single state of affairs but for many possible states of affairs. Therefore, sentences containing non-referring descriptions will still be meaningful. This aspect is derived from the logicians' definition of intension as a total function from possible worlds to entities (*i.e.*, the referents at that index). Since the number of possible worlds and of intensions is infinite, such functions are not practical for NLP systems. However, they do provide a base on which to build.
- (2) Conceptual entities (*i.e.*, intensions) are well-defined objects in our NLP system. They can be defined, referred to, and reasoned with. Both the meaning representation and knowledge representation scheme must therefore have facilities for representing intensional objects.

An *intension* will be defined as a formal mechanism to encode, in a well-defined way, the *concept* of an entity. In AI, it is useful to think of this intension as a set of salient properties sufficient to distinguish a member of its extension from any other potential referents. The content of an intension is not fixed but will vary with time and the individual. The complementary notion, that of an *extension* evaluated relative to the current context, will be identified with the entity individuated at a particular time and place. Presumably, such extensions can designate either objects in the actual (perceived) world, or imaginary objects, or states of mind⁵.

1.2. Intensions for Natural Language Processing

The logicians' model-theoretic account of intensionality is useful in what linguists term *opaque* contexts⁶. Such contexts introduce semantic ambiguities so that both a "transparent" and an "opaque" reading of the matrix clause may be possible. On the opaque reading, constituents embedded within the opaque context exhibit different semantic characteristics than they would if considered independent of the construct.

Opaque contexts are most often characterized by the failure of two standard inference rules. The first is the inability to substitute referentially equivalent expressions within an opaque context as was seen in the earlier example (1-1,1-3) of *belief*-contexts. This property is called "substitution of equals". In transparent contexts such substitution is usually permitted. Consider:

- (1-4) *The man with the martini* likes to fly.
- (1-5) *Marc Garneau* is *the man with the martini*.
- (1-6) Therefore, *Marc Garneau* likes to fly.

Statement (1-6) can be validly inferred from (1-4) and (1-5) since the phrases "Marc Garneau" and "the man with the martini" in this context refer to the same person.

⁵ Hence the extension of a phrase may turn out to be an intension.

⁶ Most of the following examples are taken from (Fodor 1979). This is a 1970 PhD thesis (published in 1979) which presents a very complete and detailed description of opacity in its various guises.

The second inference rule which fails is that of "existential generalization" as the following sentences illustrate

- (1-7) Mary *seeks* a unicorn.
- (1-8) Mary *found* a unicorn.
- (1-9) Therefore, there exists a unicorn.

From sentence (1-7) we cannot infer the truth of sentence (1-9). For the transparent case, (1-8), however such an inference to (1-9) normally is valid.

Substitution of equals and existential generalization will usually fail or succeed together but, as Fodor (1979) stresses, this need not always be the case. She argues that the two are very independent and for this reason defines two types of opacity depending on which inference rule fails. For instance from,

- (1-10) John *recognized* his long lost daughter.

one cannot conclude that

- (1-11) John *recognized* the waitress at Mac's.

even if

- (1-12) John's long lost daughter is the waitress at Mac's.

(i.e., substitution of equals fails). However, on any reading of (1-10) there is a specific person being recognized (i.e., existential generalization succeeds). Alternatively, existential generalization may fail where substitution of equals succeeds. For example, from

- (1-13) Mary *is looking for* a unicorn.

one cannot infer that

- (1-14) There exists a unicorn.

(i.e., existential generalization fails). However, from (1-13) and

- (1-15) A unicorn is a magical thing.

one can conclude

- (1-16) Mary *is looking for* a magical thing.

(i.e., substitution of equals succeeds). Further consideration will be given to such distinctions in Chapter 2 when a fuller semantic account is presented.

1.3. Purpose and Scope

The problems posed by opaque contexts can be summarized under four central questions.

- (1) How should opacity be recognized, classified, and represented?
It is the task primarily of linguists to characterize the syntactic and semantic regularities of opaque contexts, to determine how such phrases can be ambiguous, and to decide what is the best logical form (relevant to our needs) for each possible reading. These aspects are undeniably relevant for NLP systems with the added constraint that the resulting specifications be "efficiently computable".
- (2) How are inferences affected by opacity?
The two most cited inferences which fail on opaque readings are substitution of equals and existential generalization. However, it turns out that most other inference rules are also invalidated by opaque contexts (or at best their application becomes rather dubious). Such questions of logical entailment are the major concern of logicians (and to a lesser extent philosophers). Again these notions are of interest to builders of AI-NLP

systems in order to know what information can be validly inferred from an utterance and to decide what information must be explicitly represented either in the meaning representation or in the knowledge representation system.

- (3) Given that we can distinguish the two readings, what information is the speaker intending to convey on a transparent reading and on an opaque reading?
This question underlies the issues of both (1) and (2) but goes a step further by asking about the speaker's intent. More specifically, philosophers ask what information the speaker is trying to convey, what assumptions about the listener and about the world he is presupposing, and what the understander perceives. These sorts of questions are very dependent on the particular individuals involved, the lexical items chosen and, importantly, the context of the utterance. Such factors are significant for AI researchers interested in speech acts and more general pragmatics.
- (4) How should opacity be handled in AI-NLP systems?
The concerns of (1), (2), and (3), while relevant to AI and NLP, have arisen predominantly in the domains of linguistics, logic, and philosophy. The fourth question posed is influenced by the same considerations but is unique to AI. First, computational linguists must decide how to recognize opaque contexts during the parse and how to represent the possible readings in the meaning representation (MR). Given a suitable MR, the second task is to decide how to interpret this representation in a knowledge base and determine what inferences are valid⁷. Our particular concern is not with general knowledge representation (KR), but rather with its ability to model linguistic aspects of intension⁸. By examining some exemplary systems, we can state specifications which a suitable KR scheme must possess to accomplish this.

Any system that is to concern itself with the knowing and believing of information must eventually deal with intensionality. Conventional KR systems do not deal with intensionality at all — not because designers are unaware of it but rather because of the difficulties it poses. Several more recent knowledge representation schemes do attempt to handle intensions explicitly (*e.g.*, Propositional Semantic Nets (Maida & Shapiro 1982), KL-One (Brachman 1979), KRL (Bobrow and Winograd 1977), EFR (Nishida 1983)). Even though such KR schema can provide insights into the representation of intensions, they do not explicitly examine their role in NLP systems. That is, they do not address the problem of parsing opaque constructs, of determining the relevant ambiguities, nor of how to disambiguate the resulting meaning representation. It is these aspects that I wish to consider. Three major themes dominate the scope of the thesis:

- (1) motivating the need for and defining a semantic representation to handle linguistic opacity,
- (2) demonstrating how such a semantics can be made practical for AI systems by combining logic programming and Montague-style semantics,
- (3) discussing how such a semantic representation can successfully be integrated with knowledge-based systems.

My goal is not to "solve" the problems posed by opaque contexts, but rather to propose a concise set of initial specifications for the correct handling of them within a computational framework.

⁷ We also must consider how these inferences are to be computed which is another hard problem in non-first order logics.

⁸ We do not advocate any particular KR representation. In fact one of our motivations for having an intermediate meaning representation is to allow for flexibility in our choice of KR scheme.

1.4. Thesis Overview

In Chapter 2 we identify the syntactic constructs in which opacity may arise, and present an outline of the semantic considerations relevant to these contexts. Chapter 3 introduces the idea of *descriptors* and other terminology for discussing in fuller detail the semantic features of opaque contexts and suggests a classification scheme based on these distinctions.

Chapter 4 defines the requirements and operators of a meaning representation for opaque constructs, while Chapter 5 focuses on the substitution of equals, stressing the wide range of variability and the important role of descriptors for a proper analysis. A minimal set of requirements for substitution is proposed. Finally, in Chapter 6 the chosen meaning representation is justified by describing its use in a logic-programming based NLP system and how it could be effectively integrated with standard AI models for knowledge representation. We then summarize the major points of our research and contemplate directions for future work.

CHAPTER 2

Linguistic Motivations for Intensions

The ambiguities introduced by opaque contexts are essentially semantic and/or pragmatic. Nevertheless it is important that a NLP system be capable of recognizing the syntactic constructs in which they may arise. While there does appear to be some structural regularity in the occurrence of opaque contexts, they are not identifiable from syntactic properties alone. A general rule of thumb is that all noun phrases (NPs) which appear in complement clauses can be ambiguous with respect to opacity. In addition NPs in non-complement sentences containing modal verbs, intensional transitive verbs, intensional adjectives, and sentential adverbs can be ambiguous.

2.1. Syntactic Classification

The following attempts to categorize, according to syntactic form, those constructs in which ambiguity due to opacity can occur. Some of the defined categories turn out to be semantically equivalent, while others will require further distinctions to differentiate among the members.

(1) Sentential Complement Clauses:

Syntactically these take the form of "X that ϕ " where ϕ is a sentential complement clause and X is an expression of one of the following types.

(i) *It*-Cleft Phrases:

It is possible, It may be, It is necessary, It is apparent, . . .

(Exceptions: *It is true that, It is not the case that.*)

(2-1) *It is possible that Mary's murderer is insane.*

These are characterized as sentence embedding constructs with an impersonal subject. The speaker is expressing a proposition contingent on some possible state(s) of affairs.

(ii) Propositional Attitudes:

knows, believes, doubts, supposes, imagines, . . . (Epistemic)

wishes, hopes, intends, proved, surprised, . . . (Emotive)

heard, saw, recognized, . . . (Perceptive)

(2-2) *Terry believes that Mary's murderer is insane.*

Sentences in this subclass express a relation between an individual and an embedded sentence. (Often the "that" is implicit in the surface syntax but the ϕ constituent is still a full-fledged sentential complement.)

(iii) Reported Speech:

said, reported, remarked, announced, . . .

(2-3) *John announced that he was to become a father.*

The speaker is conveying the form and/or content of another's verbal behaviour, either directly or indirectly. For direct quotation the "that" is replaced by quotation marks as in

(2-4) *Stephen said, "Grapes are green".*

The surface syntax of classes (2) to (7) is in a non-complement clause form.

However, the majority have natural paraphrases which do exhibit complement structure.

(2) Modal Verbs:

These auxiliary verbs act as modal operators on a sentence to express a contingent proposition whose conditions may or may not be explicitly expressed.

(i) Standard Modals:

will/would, can/could, may/might, shall/should, must/have to, . . .

(2-5) Philip *might* marry a beautiful woman.

(ii) Counterfactuals:

(2-6) If everybody *were* to mind their own business,
then the world *would* go round a deal faster than it does.¹

Counterfactuals take the form of "If it *were* the case that ϕ , then it *would* be the case that ω ", where ϕ and ω are sentential constituents.

(3) Sentential Adverbs of Attitude:

necessarily, probably, possibly, perhaps, apparently, conceivably, obviously, understandably, . . .

(2-7) *Probably* the man in the dark trench coat ate your bagel.

These adverbs, used to modify declarative sentences, convey the speaker's comment on the content of what he is saying².

(4) Intensional Transitive Verbs:

Intensional verbs can be subclassified according to the types of objects they can take.

(i) Those taking *either* infinitive complements or simple NPs as objects:

wants, hopes (for), expects, needs, . . .

(2-8) John *needs* a yacht. (Simple NP)

(2-9) John *needs* to have a yacht. (Infinitive Form)

(2-10) *John *needs that* John has a yacht³. (Paraphrase)

Both the infinitive form and the simple NP form can be used to express the same proposition. Most can be awkwardly paraphrased into a sentential complement form by inserting "that" with the same subject as the parent clause, as (2-10) illustrates.

(ii) Performatives:

ordered, commanded, . . .

(2-11) John *ordered* a triple-decker sandwich. (Transitive)

(2-12) *John *ordered that* John be given a hamburger. (Paraphrase)

(2-13) John *ordered Mary* to shoot a squirrel. (Ditransitive with Infinitive)

(2-14) ?John *ordered that* Mary shoot a squirrel. (Paraphrase)

Verbs of this subclass can be transitive and take a simple NP object, or ditransitive and take an infinitive complement as object. In contrast to group (i), whenever performatives take infinitive complements they must also have an indirect object. To express a particular proposition only one of the two forms is acceptable. The infinitive form is

¹ Lewis Carroll, *Alice's Adventures In Wonderland*, Bantam Classics Edition, 1981, p. 43.

² Contrast sentential adverbs of "style" such as *seriously, personally, and loudly* which comment on the form or force of what is said rather than the content. (Quirk 1984)

³ Following the convention in linguistics, "?" marks semantically dubious phrase, and "*" marks semantically unacceptable phrases.

more amenable to paraphrase than is the simple NP form.

(iii) Those taking **only** simple NPs as objects:

looks for, seeks, worships, conceives of, hunts, recognize, looks like, . . .

(2-15) Bill is *looking for* a hammer.

For this subclass it is difficult to find semantically equivalent paraphrases that exhibit sentential complement structure. (e.g., Does replacing "seek" by "try to find" capture the same meaning? What is a reasonable paraphrase for "worships"?)

(5) Intensional Adjectives and Adverbs:

*famous, former, alleged, good, competent, fake, ostensible, sincere, typical, . . .*⁴

(2-16) John is a *former* senator.

(2-17) John was *formerly* a senator.

Intensional modifiers serve to restrict the set of individuals referred to by the associated common noun. The meanings of such adjectives or adverbs are highly context dependent.

(6) Copula Verbs:

be, become, resemble, equal, represent, . . .

(2-18) The Chief of Police *became* a member of the City Council.

Such verbs can be used in *attributive* clauses to predicate an entity, or in *equative* clauses to assert identity between two descriptions. Since these verbs can relate intensional as well as extensional objects, they are often used to establish criteria for substitution and inferencing in opaque contexts.

(7) Other Candidates:

(i) Intensional Prepositions:

(2-19) Don ordered a book *about* a butterfly.

The use of "of" and "about" can produce opaque ambiguities. In (2-19) it is not clear whether the object NP refers to a *particular* butterfly or just *some* butterfly.

(2-20) John is *talking about* a unicorn.

Some combinations of verbs plus particles create opaque contexts but these are actually intensional transitive verbs of which the preposition is a part. This is the case in (2-20). Others of this type include:

speak about/of, silent about, tell about/of, talk about, . . .

(ii) Subordinate Conjunctions:

Conjunctions expressing causality such as *because, since, while* do not permit substitution of sentential clauses. For example, in a situation where both "the grass is brown" and "John is an engineer" are true, this fact does not give us licence to intersubstitute them. Thus, we cannot infer (2-22) from (2-21).

(2-21) *Because* [the grass is brown], the valleys don't appear very lush.

(2-22) *Because* [John is an engineer], the valleys don't appear very lush.

The expectation that the clauses can be substituted stems from using a truth conditional analysis. Nonetheless the intensions of the two conjoined clauses must somehow be related causally. It is not clear how this is to be determined. It seems we need an internal structure for concepts along which they can be compared.

(iii) Creative Verbs in Progressive Tense:

⁴ See Ali (1985) for a fuller list and discussion.

(2-23) *I am writing a thesis.*

The object NP of a creative verb is not presumed to exist; it may partially exist and an object of the corresponding type should be realizable. It is a question of existence and not specificity. The referred to entity will come into existence when the process described by the verb is completed.

(iv) Interrogative Contexts:

(2-24) *Who is Mary's boyfriend?*

(i) The world's greatest pianist.

(ii) That man in the green chair.

Interrogative contexts (including yes/no questions, term questions, WH-questions) are similar to modal contexts in that they consider possible rather than actual states of affairs.

Table 1 summarizes the syntactic possibilities for opaque contexts. Almost any context can be made opaque by altering the normal stress and intonation when speaking, or adding surrounding quotes to an embedded phrase as in

(2-25) Jack regrets that his 'silly little sister' is now in jail.

These types of opaque contexts are actually easier to handle since they provide explicit information as to what the intended reading is.

2.2. Semantic Considerations for Opaque Contexts

Syntax alone is insufficient to characterize opaque contexts. A given sentence can be structurally unambiguous but have several semantic representations because of the presence of opacity. We must appeal to semantic and pragmatic features to resolve the ambiguity. Earlier we noted that opaque contexts affect the logical entailments and presuppositions of the sentence being asserted. We will now look at these in more detail.

2.2.1. Inference Failures

In formal semantic theory, a sentence p *logically entails* another sentence q if and only if whenever the conditions of p are satisfied, the conditions of q are also satisfied. In natural language the hearer is not normally expected to be aware of the infinite number of logical entailments of a sentence⁵ and certainly the speaker does not intend to convey all of these entailments by his utterance. A more workable notion of logical entailment comes from *relevance* logic (Belnap and Anderson 1975) and restricts the implicatures to those derived solely from the semantic content of the sentence which the speaker would expect the hearer to infer (i.e., are relevant to the utterance)⁶.

Inference rules state explicit cases when logical entailment will hold between two propositions, and thus indicate what further information is being conveyed by a linguistic phrase. Certain rules of logical inference fail to apply within opaque contexts:

2.2.1.1. Existential Generalization

In logic, the rule of *existential generalization* states that for any constant, a , $P(a) \rightarrow \exists x: P(x)$, where P is some unary predicate. This correctly predicts that from

⁵ Such an assumption, termed *logical omniscience*, is clearly unrealistic (Levesque 1983).

⁶ This type of semantic implicature should be distinguished from pragmatic implicatures which are derived from both the content of what has been said and some specific assumptions about the co-operative nature of verbal interaction. The latter conform to the *conversational implicatures* of (Grice 1975) and involve many context-dependent and pragmatic factors. We do not consider those here.

SYNTACTIC CLASSIFICATION

Sentential Complement Clauses

It-Cleft Phrases	<i>It could be that</i> the Big Apple is an island.
Propositional Attitudes	Pat <i>believes that</i> a clown wants to meet Jim's wife.
Reported Speech	John <i>announced that</i> Bill's wife was pregnant.

Modal Verbs

Standard Modals	Derek <i>would like</i> to marry a girl his parents disapprove of.
Counterfactuals	If McDonalds <i>were</i> granted the concession downstairs, then the students <i>would</i> starve.

Sentential Adverbs of Attitude

Certainly the Prime Minister will do the best for his party.

Intensional Transitive Verbs

Complement or Simple NPs	John <i>wants</i> the fastest car in the world.
Performatives	Colonel Remy <i>commanded</i> the entire French Army.
Simple NPs	Craig <i>worships</i> a dragon with three heads.

Intensional Adjectives and Adverbs

Julianne bought a *fake* gun.

Intensional Prepositions

Do you have a poster *of* a unicorn?

Copula Verbs

Jerry *resembles* a famous politician.

Subordinate Conjunctions

Since the groceries are bought,
Rob can cook dinner.

Creative Verbs

Sue *is making* a telephone call.

Interrogative Contexts

Who is the winner of the book prize?

TABLE 1

(2-26) A man walked into the room.

(2-27) There exists a man who walked into the room.

The failure of this rule in opaque contexts, almost exclusively affecting noun phrases, can be construed in two ways. The first stems from a strict logical interpretation of the rule and the second from a more linguistic interpretation.

(i) As a Lack of Requirement that the object NP refers to anything in the Actual World:

A noun phrase may describe an object that exists in the "actual" world but within the scope of an opaque construct it need not. Two readings are possible depending on what presumptions we attribute to the subject. For example, consider the sentences:

(2-28) John *believes that* a man walked into the room.

(2-29) There is a man who walked into the room.

(2-30) John *believes that* there exists a man who walked into the room.

From sentence (2-28) we cannot infer (2-29), while on the transparent version (2-26) we could. On an opaque reading of sentence (2-28), the standard inference of existential generalization fails. Even though the speaker need not be committed to the existence of the object, such a commitment is being attributed to the subject so (2-30) can be concluded. This phenomenon will be termed ambiguity due to *realizability*.

(ii) As a Lack of a Specific Referent for an object NP:

We distinguish between specific and non-specific readings for indefinite descriptions, and between referential and attributive readings for definite descriptions. On a specific reading the NP identifies one extensional object as being referred to; whereas on a non-specific reading no particular entity is individuated, the referent is intensional. Both readings may be plausible and so there is ambiguity due to the *intensionality of the referent*.

(2-31) John *wants to* talk to [a man in a brown three-piece suit].

On the opaque reading of (2-31), John wants to talk to some person fitting the description "man in a brown three piece suit" but on the transparent reading there is a particular person whom John has in mind and he has used the phrase "man in a brown three-piece suit" to refer to him even though he could perhaps have used some other phrase. Intuitively, the opaque reading of an NP allows one to believe something without there having to be anything one believes it *of*, and this is why existential generalization fails on opaque contexts.

2.2.1.2. Substitution of Equivalents

Standard logics prescribe that whenever $(a=b)$ and $P(a)$ are known to be true, then $P(b)$ must also be true. We would expect then that phrases referring to the same individual could be intersubstituted. Within opaque contexts however this is not generally the case. The failure of this rule can again be interpreted in either of two ways, one logically and the other linguistically.

(i) Failure to Substitute Referential Equivalents:

The following contain constituents which refer to the same value and yet can not be intersubstituted.

(2-32) Ralph *thinks that* [the number of planets] is seven.

(2-33) [The number of planets] is [nine].

(2-34) Therefore, Ralph *thinks that* [nine] is seven.

From the logicians' perspective, the phrases "the number of planets" and "nine" from (2-33) are co-referential; both refer to the number nine. However, they differ in conceptual meaning and the expected syllogism fails.

(ii) Opacity with respect to Descriptive Content:

A related phenomenon but more linguistically based, Fodor (1979) calls *opacity with respect to descriptive content*. This concerns who is taken as accepting the descriptive content of the phrase in question. Good examples are those involving self-contradictory clauses. Even though the sentence

(2-35) [Your dead uncle] is alive.

is contradictory (and even peculiar) on its own, it becomes semantically ambiguous in *belief*-contexts such as

(2-36) Heather *believes that* [your dead uncle] is alive.

On a transparent reading of (2-36) only the speaker is taken to subscribe to the description

"your dead uncle". The speaker is attributing a mistaken belief to Heather. However, on an opaque reading, both the speaker and Heather accept the description "your dead uncle" and so Heather is attributed a contradictory belief (perhaps a belief in reincarnation). In general, the speaker is always responsible for the description; the ambiguity is a matter of whether or not the description is also ascribed to the agent(s) of the opaque verb(s). Attributing the description to both speaker and agent(s) is considered the opaque reading, and attributing it only to the speaker considered transparent. Before any substitution is allowed, we must consider to whom the descriptive content should be attributed. This is discussed in section 5.4.3.

Jackendoff (1983) nicely defines this so-called "Opacity Principle" for *belief*-contexts:

Suppose that a person P holds a belief B about some entity E. In describing B, a speaker may describe E either

- a) in terms of P's internal representation of E [opaque description], or
- b) in terms of a representation that adequately identifies E for the speaker and hearer, though not necessarily for P [transparent description]. (Jackendoff 1983 : 214)

The intuition here is that on the opaque reading one can believe something if it is described in one way without necessarily believing it if it is described in another way. Interestingly, this Opacity Principle does not reflect a rule of logic, but rather a feature of the semantics of the particular natural language. Jackendoff (1983 : 214) points out that some languages (*e.g.*, Malagasy, Fering) have grammatical markers to distinguish this type of descriptive ambiguity.

2.2.1.3. Other Logical Inferences

The principle of *universal generalization*, which states that for any constant, a , $((\forall x P(x)) \rightarrow P(a))$, fails in some opaque contexts. Thus,

(2-37) John *said* that [everyone] loves popcorn.

does not imply that

(2-38) John *said* that [Joe Clark] loves popcorn.

Also, the *simplification of conjuncts*, which permits the reduction of $P(a \text{ AND } b)$ to $P(a)$ AND $P(b)$, does not apply within opaque contexts. This question can arise if, as some linguists advocate, restrictive relative clauses are represented as conjuncts. For example, the sentence

(2-39) John is upset *because* [Bill and Jane] got married.

should not allow the hearer to conclude that

John is upset *because* [Bill got married] AND

John is upset *because* [Jane got married]

2.2.2. Presupposition

Closely tied in with logical entailment is the notion of semantic presupposition.

A sentence p is *semantically presupposed* by another sentence q if and only if

- (a) in all situations where q is true, p is true, and
- (b) in all situations where q is false, p is true, and
- (c) if p is ever false, q is neither true nor false.

(Levinson 1983 : 175)

A sentence will only be considered meaningful if all the conditions specified by its presuppositions are satisfied.⁷ Certain constructs have been isolated by linguists as sources of presupposition. Those of interest to us because of their connection with opacity include⁸:

- (i) **Definite Descriptions:** In transparent contexts the existence and uniqueness of entities referred to by a definite description is presupposed, but need not be in opaque contexts. This contrasts with indefinite descriptions which do not *presuppose* existence but may *entail* existence.

- (ii) **Factive Verbs:** Factive verbs express an attitude towards a proposition whose content is taken to be factual. The sentence

(2-40) *Martha regrets that she drank John's home brew.*

presupposes the truth of

(2-41) *Martha drank John's home brew.*

- (iii) **Counterfactuals:** The counterfactual sentence

(2-42) *If Hannibal had only had 12 more elephants,*

then the Romance languages would/would not this day exist.

presupposes the negation of the antecedent, namely that "Hannibal did *not* have 12 more elephants". With counterfactuals of the general form "If ϕ , then ω ", the opposite proposition to that expressed in the *if*-clause is presupposed.

Presuppositions, while affected by some types of opaque constructs, are not a characteristic feature of such constructs. As well it is difficult to give a systematic account of them without introducing a large pragmatic component unrelated to opacity.

2.2.3. Deletion of Modifying Phrases

In opaque contexts phrases post-modifying nouns cannot be deleted even though they can be deleted in transparent contexts. Consider:

(2-43) Bill met a pretty girl [that didn't like him].

(2-44) Bill met a pretty girl.

(2-45) Bill *would hate* there to be a pretty girl [that didn't like him].

(2-46) Bill *would hate* there to be a pretty girl.

It is valid in the transparent context of (2-43) to infer (2-44), but the opaque context (2-45) does not entail (2-46). (Fodor 1980 : 188)

2.3. Noun Phrases, the prime target for ambiguity

Embedded noun phrases are a prime source of ambiguity in opaque contexts responsible for subtle but significant distinctions in meaning. There are several subclasses of NPs, each affected somewhat differently by opacity. Traditionally these are divided into two major categories based on the syntactical properties of the NP: definite descriptions and indefinite descriptions (see Appendix A for further examples and details).

- (1) **Definite Descriptions:**

⁷ This implies the need for a three-valued logic consisting of TRUE, FALSE, and UNDETERMINED. If truth values are to be significant, this might be a reasonable idea anyway in light of the difficulties handling incomplete knowledge in a two-valued logic.

⁸ Karttunen (n.d.) as reported in Levinson (1983).

Definite descriptions uniquely identify one individual or one set of individuals. They are used either *referentially* to pick out a certain individual, or *attributively* to specify the properties sufficient to pick out some individual. In both cases, the speaker presupposes a prior awareness of the referent. This can be established linguistically from the preceding discourse, or non-linguistically from the discourse situation.

Noun phrases consisting solely of a proper noun, personal, demonstrative or possessive pronoun, and those composed of a common noun preceded by a demonstrative article are all considered definite descriptions and are always read referentially. Common nouns preceded by a definite or possessive article are also definite descriptions but can be read either referentially or attributively and thus are affected by opacity.

(2) Indefinite Descriptions:

Instead of individuating *particular* members, indefinite descriptions serve to identify some object or set of objects as belonging to a certain class. In this respect indefinite descriptions act like specifications for set membership. Such noun phrases can be read as either *specific*, to indicate that the speaker has a particular individual in mind, or *non-specific*, when any member of the class will do. This is somewhat different than the referential/attributive distinction of definite descriptions in that no prior awareness of an individual of that sort is presumed.

Question and quantified pronouns, mass nouns, and generic noun phrases are classes of indefinite descriptions which are always read non-specifically. It is common nouns preceded by an indefinite article and indefinite quantified NPs which must be disambiguated in opaque contexts since they can be read either specifically or non-specifically.

Definite and indefinite descriptions can also be read *generically* in certain contexts. Thus, in

(2-47) A Honda Prelude comes equipped with a sun roof.

(2-48) The brontosaurus was an herbivorous dinosaur.

the descriptions in the subject position refer to a typical member of a class rather than an individual.

From a syntactic point of view this categorization may suffice, but when ascertaining their semantics, especially the intensionality of the referent, further refinements are needed. Mallery (1985) suggests that a third type, *class* descriptions, should be defined to include NPs describing sets or classes of objects. Examples would be NPs quantified by "all", "any", "every", and "no". Even though such a distinction seems useful, in order to avoid extra complexity I restrict myself to the use of definite and indefinite for classifying descriptions. (In fact, I only consider singular NPs quantified by "the" and "a". The distinctions for opacity should be equally applicable to other types of determiners.)

2.4. Factors Influencing Ambiguity

We restrict our attention to the failure of the first two logical inference rules and their linguistic counterparts, namely *existential generalization* and *substitution of equals*. This will necessitate distinguishing among three types of ambiguity:

- (i) **Realizability:** A noun phrase may or may not be interpreted as having a referent that currently exists in (listener's model of) the world.
- (ii) **Intensionality of the Referent:** A noun phrase can refer to an extension or an intension depending on whether it describes a *specific* entity or not.
- (iii) **Descriptive Content:** The descriptive content of a noun phrase can be attributed to agents introduced as subjects of intensional constructs as well as to the speaker.

Several factors will influence the type of ambiguity present and must be taken into account during the parse. These factors are discussed in the following sections.

2.4.1. Definiteness of Noun Phrases

In non-opaque contexts, definite descriptions are generally referential (unless marked semantically or syntactically as generic)⁹. However in opaque contexts, definite descriptions can be given referential, attributive, or generic readings.

(2-49) John *thinks that* [the lion in my backyard] is a fierce beast. (Referential)

(2-50) John *thinks that* [the lion in Africa] is a fierce beast. (Generic)

(2-51) John *thinks that* [the lion] is a fierce beast. (Ambiguous)

Similar comments apply to indefinite descriptions. The choice of reading will affect the intentionality of the referent.

2.4.2. Degree of Description of Noun Phrases

Degree of description refers to the amount of detail included in a noun phrase, perhaps measured by the number and type of modifiers. This factor reflects the size of the set of entities being described and can affect the specificity of the noun phrase. The more descriptive a given noun phrase is, the more likely it is to be interpreted as specific rather than non-specific.

2.4.3. Verb Tense and Aspect

The choice of verb tense and aspect influences the modality of a construct and the existence presumptions of a referring description. Future and conditional tenses and the progressive aspect are most relevant to opacity. Compare the existence presuppositions of the following three sentences:

(2-52) She *will* write a paper. (future - no existence)

(2-53) She *is writing* a paper. (progressive - partial existence)

(2-54) She *wrote* a paper. (simple past - complete existence)

2.4.4. Verb Type

Stative verbs express states of perception and cognition and changes among them. In contrast, *dynamic* verbs describe either physical or mental events and actions¹⁰. With verbs of physical action the subject and object noun phrases are usually given a specific reading unless the verb is in a future or modal tense. The meaning of stative verbs on the other hand are much more context dependent and therefore more susceptible to ambiguities of opacity.

2.4.5. Surrounding Linguistic Context

The amount of opaque ambiguity in a sentence depends on how deeply a constituent is embedded within opaque constructs. A given noun phrase can be ambiguous with respect to each clause it is embedded in. In the sentence,

(2-55) Mary *believes that* [Jack *wants* [a yacht like Bill's]].

the noun phrase "a yacht like Bill's" can be considered relative to just the *want*-clause, or to both the *want*-clause and the *belief*-clause. For each, there are two types of ambiguity to

⁹ Prior context can cause definite descriptions to be read non-specifically but, for the time being, we consider only sentences in independent of context (see Strzalkowski and Cercone 1985).

¹⁰ An ACTION is an EVENT controlled by an AGENT.

consider (*i.e.*, descriptive content, intensionality of the referent) so that this sentence has nine potential readings each varying slightly in meaning.

The above are immediate *intrasentential* features which influence the amount of ambiguity we need to represent. For *disambiguation*, the most vital information is usually pragmatic (*i.e.*, discourse environment, intent of the speaker, and non-linguistic assumptions) rather than semantic.

A final point concerns the inheritance of transparency or opacity from subconstituents. For determining the *intensionality of the referent*, it is opacity which is inherited by a parent from its subconstituents. If a subconstituent is given an intensional reading then the parent constituent must also be given an intensional reading. This is dictated by the inheritance of specificity. For instance in (2-57) if "a yacht like Bill's" is read non-specifically relative to the *want*-clause, then it must also be non-specific with respect to the *belief*-clause.¹¹ In contrast, for ambiguity due to *descriptive content* it is transparency that is inherited. If the subconstituent is read transparently then the parent phrase must also be read transparently. So, if "a yacht like Bill's" is *not* attributed to Jack at the *want*-clause level then it can not be attributed to Jack at the *belief*-clause level.

¹¹ It is conceivable that the NP could be read specifically with respect to the *want*-clause and non-specifically with respect to the *belief*-clause. But I'm not convinced this is needed and, if it is, how it could be recognized.

CHAPTER 3

Semantic Classification

Our aim is to make the criteria for applying inference rules in opaque contexts more concise and thus amenable to computation in an NLP system. To facilitate subsequent discussion of our semantic classification, we need first to establish some terminology and distinctions upon which our classification is based. The counterpart of a concept in our NLP system is an *intension*, a finite representation of those properties which are able to characterize an instance of the corresponding concept in a given state of affairs. The properties that actually constitute the known intension will be different for each person and, at any given time, the cognitive agent (*i.e.*, the NLP system) can only have incomplete knowledge of this set. Thus an intension consists of a dynamic set of properties that can be expanded and modified with experience (*cf.* KL-One).

An idea central to our treatment of opacity is that of a *descriptor*, defined to be any non-empty subset of an intension (in practice, often identical to the intension). A descriptor provides access either to the intension of which it is a part or to its extension, eliminating the need to explicitly list all the known properties; only properties relevant to the discourse situation are mentioned.

3.1. The Use and Evaluation of Descriptors

Noun phrases are used to describe objects. In transparent contexts the entity referred to is extensional; it is presumed to correspond to something being introduced or already in the discourse environment. In opaque contexts the descriptive role of NPs is accentuated and interpretation of these descriptions becomes an important issue, one that has been overlooked by most AI-NLP designers. The comments about descriptions are relevant to several types of syntactic constituents (Fodor 1979 : 282). We focus on noun phrases (in opaque contexts termed *descriptors*) and sentences (termed *propositions* in opaque contexts) which provide ample food for thought.

Descriptors and propositions may vary along several dimensions as illustrated in figures 1 and 2. The set of dimensions is similar in both cases, although the precise interpretation of each may be slightly different. The dimensions include:

WHAT: A NP descriptor can be used to access either the intension or the extension of an entity. To decide this will involve determining the specificity of the NP description.¹ The choice for sentential constituents is whether the proposition refers to the set of conditions it embodies or to their satisfiability in the current state of affairs (*i.e.*, its truth value).²

WHEN: A second aspect of reference concerns when the description is intended to apply. NP descriptors may always apply (*e.g.*, proper names), may apply at some particular future time, or may not apply at all. Likewise, propositions can describe conditions that are

¹ The usual existential presuppositions for indefinite and definite descriptions still hold: definite descriptions assume some prior awareness of a unique referent, while indefinite usually do not.

² The truth value is either *true* (*i.e.*, satisfied) or *false* (*i.e.*, not satisfied) in a state of affairs. If a value is not specified, it is assumed to be unknown.

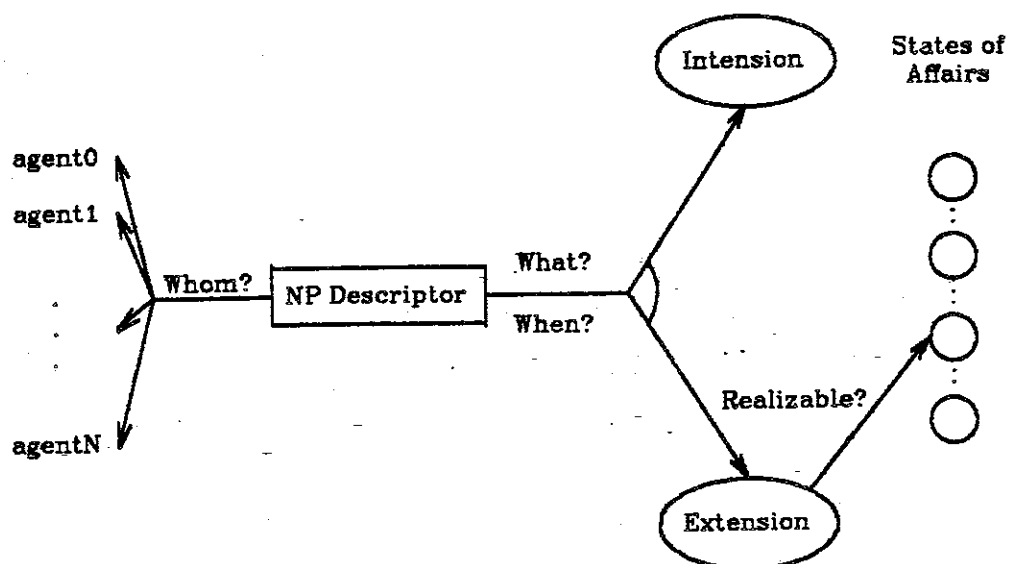


Figure 1. Dimensions of Descriptors

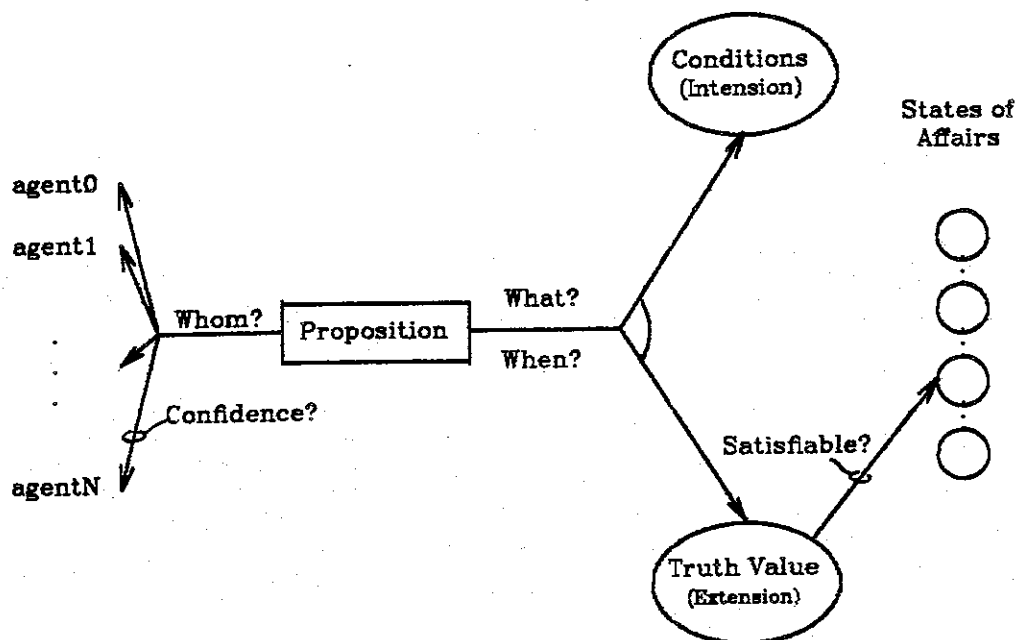


Figure 2. Dimensions of Propositions

already satisfied, satisfied only in hypothetical states of affairs, or satisfied at some future unknown time. Much work done in temporal logics and reasoning has focused on this dimension.

WHOM: Descriptors and propositions are always attributed to the speaker, unless the speaker is quoting.³ In opaque contexts containing ambiguity of descriptive content, it is essential to articulate further whether descriptors and propositions are to be attributed to other agents as well. This prevents invalid inferences from being drawn and enables correct user models to be constructed in the knowledge base (KB).

REALIZABILITY: Extensional referents for descriptors only "exist" in certain states of affairs. Likewise, propositions have truth values only in certain states of affairs. The relevant dimensions are marked in the diagrams as *Realizable* (for descriptors) and *Satisfiable* (for propositions). Since states of affairs are just symbolic representations in the KB and have no necessary correlation with the real world, it must be decided which correspond to the "real" (perceived) world and which are just imaginary. Such questions of real-world existence are often confused with issues of specificity. However, this is a separate issue that has to do with KB modelling and constructing agents' belief sets. It is up to the host KR system to mark those entities and conditions which are believed to correspond to the real world. For cases where "real" existence is significant in opaque contexts a *believed-to-exist* predicate must be used.

CONFIDENCE: An utterance may express a speaker's degree of conviction towards an embedded proposition. This reflects how certain the speaker (or another agent) is that the proposition is attainable from the current state of affairs, which affects substitution of equals as described in chapter 5. This confidence level is not as straightforward as one might hope, since it involves reasoning about hypothetical states of affairs (there are many such states but only a few are relevant), and the confidence expressed can be towards not only the proposition but also towards the inferential process required to determine the proposition (*e.g.*, luckily, obviously).

3.1.1. Referents for NP Descriptors

Syntactically, a quantified NP descriptor is classified as definite or indefinite (or class) according to the type of its determiner. Such descriptors are then read as being *specific* or *non-specific* or *generic* for indefinite descriptions, and *referential*, *attributive*, or *generic* for definite descriptions. This is influenced by the definiteness of the determiner but not dictated by it. In a NLP system each of these readings must be given a representation in the KB, its "knowledge base interpretation". Standard KR schemes permit referents to be either *individuals* (*i.e.*, INSTANCES) or *concepts* (*i.e.*, CLASSES). Figure 3 summarizes a classification scheme for singular noun phrase descriptions. For a proper interpretation of non-specific readings, I suggest that *individual* concepts need to be distinguished from *universal* concepts⁴.

³ In general, if a speaker takes over an agent's description, the speaker commits himself to its accuracy. However, if the speaker is quoting, he is not committed to the accuracy (or use by himself) of the quoted fragments.

⁴ In the lingo of KR types, our term *individuals* corresponds to INSTANCES, *universal concepts* to CLASSES, and *individual concepts* to something like INSTANCE-OF slots.

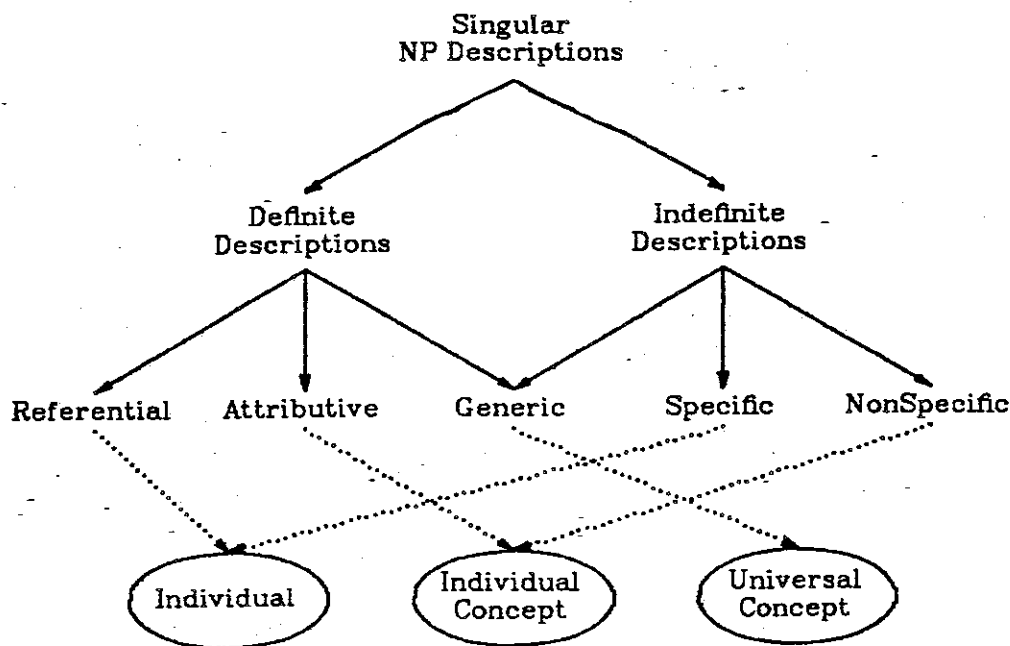


Figure 3. Referents for NP Descriptions

The crucial issue is to decide whether non-specific and generic readings of descriptions are sufficiently distinct to warrant two distinct types of conceptual node in a KB. Most often the two, generic and non-specific, are not distinguished. Consider the following set of examples:

- | | | |
|--|------------------------|------------------------|
| (3-1) <i>A dog</i> ran across the street. | : specific reading | → individual |
| (3-2) <i>A dog</i> has four legs and a tail. | : non-specific reading | → universal concept |
| (3-3) Mary wants <i>a dog</i> . | : specific reading | → individual |
| | : non-specific reading | → ?individual concept? |
| | : generic reading | → ?universal concept? |

On the specific reading of (3-3), Mary wants a *certain* dog. The interpretation of the non-specific reading is more debatable; if "a dog" refers to a universal concept it would seem that Mary wants *any* dog, while if it refers to an individual concept then Mary wants *some* dog; that is, one (unspecified) member of the class DOG. In this example, the latter seems more appropriate.

Although, the distinction between the two may seem strained in this example, it is necessary, as the next example shows.

- (3-4) *A friend of mine* is a very lucky person.

This sentence is ambiguous between generic reading (3-5) in which a universal concept is the referent, non-specific reading (3-6) in which an individual concept is the referent, and a specific reading such as (3-7) in which an individual is the referent.

(3-5) *Any* friend of mine is a very lucky person.

(3-6) *One* of my friends is a very lucky person.

(3-7) My friend *Sue* is a very lucky person.

The important point is that distinct generic and non-specific readings do exist, and that their referents will require different KB nodes if they are to be distinguished. Some may argue that such a distinction is too fine-grained and results from the particular determiner-noun combination rather than being a general phenomenon. However, *individual* concepts are especially useful for interpreting sentences like:

(3-8) No two presidents of the U.S. *ever* looked alike.

for which a *universal* concept reading of "President of the U.S." would be problematic.

3.1.2. Equivalence Among NP Descriptors

Asserting that two noun phrase descriptions are "equal" may mean any of the following:

- (1) They are co-extensional (or co-referring) descriptors: The descriptors have the same referent in a particular state of affairs (usually the current one) but not in all states of affairs.

(3-9) *The winner of the New York lottery is the owner of Pizza-Pizza.*

- (2) They are co-intensional descriptors⁵: The descriptors refer to intensions which happen to co-refer in all states of affairs.

(3-10) *The Morning Star is the Evening Star.*

- (3) They are definitionally equivalent descriptors: The two descriptors express the same *concept*; they are synonyms.

(3-11) *A woodchuck is the same as a groundhog.*

Case (3) implies (2) and (2) implies (1), but the converses of course do not hold. Previous accounts overlook the distinction between (2) and (3), treating both relations as co-intensions. This is a source of difficulty in Montague semantics where the two types are conflated. The sentence

(3-12) [the Morning Star] is [the Evening Star].

exemplifies a type (2) relation by asserting that "the Morning Star" and "the Evening Star" will always refer to the same object but it certainly does not assert that the underlying concepts are the same. In contrast,

(3-13) [Oak tag] is [bristol board].

does assert that two descriptions express the same concept, a type (3) relation.

Type (3) equivalence arises most often when synonyms are given, when terms are being completely defined as in:

(3-14) [The kiwi] is [the flightless bird native to New Zealand].

and among paraphrases like "a fruit with an orange skin" and "an orange-skinned fruit". Neither are particular to opaque contexts. When the need arises to assert definitional equivalence we use the relational symbol "*=_d*"; paraphrases will be considered as equivalent descriptors with no symbolic distinction.

⁵ Using this terminology is perhaps a bit confusing — it is based on the logicians' definition whose only criteria for equivalent intensions is that they co-refer in all contexts; they assert nothing about the underlying concepts being the same (although most people read this into it).

To clarify the distinction between co-*extensional* and co-*intensional* descriptors, consider the three NP completions below. All are presumed to refer to the current Prime Minister of Canada in some contrived context, say a spring cabinet meeting:

(3-15) The reporter *wanted* to interview

- (i) the man in the blue-striped suit,
- (ii) the Prime Minister of Canada,
- (iii) the leader of the party with the most members in the Commons.

We can imagine situations in which completions (i) and (ii) are co-extensional descriptors, but do not expect them *always* to have the same referent. However (ii) and (iii) would be considered co-intensional descriptors over most reasonable states of affairs.⁶

3.1.3. Labels, Special-Purpose Descriptors

Definite and indefinite descriptions identify objects by isolating a property that an object has. Proper nouns also identify entities but by associating names with them; this association is by definition rather than some property of the object. We interpret proper nouns as *labels*, special-purpose descriptors which are *always* referential, refer to the same entity across all states of affairs, and for which knowledge of the labelled entity is presumed. Given a label, a listener immediately knows or doesn't know to what it refers, but given a descriptor, he must deduce from the state of affairs and discourse context what is being referred to. We assume a one-to-one correspondence between labels and their denotations. That is, any denotation of an entity (*e.g.*, an instance number in a KB system, an individual constant in model theory — whatever is to represent the “real” thing) can have a unique label associated with it. By using the label, we are in essence referring to that individual⁷.

For conversational efficiency, speakers tend to use the most specific referring phrase warranted by the situation. All else equal, then, a speaker will choose a label over a descriptor. The use of descriptors is however preferred in a number of situations:

- (1) The speaker has a particular individual he wants to refer to. He may not know the label of this individual and so must choose an alternative, namely a descriptor read specifically, which is his only linguistic means of referring to that individual.
- (2) The speaker may believe that the listener is unaware of the label of an individual. Perhaps the listener doesn't know the individual at all, or he may know of the individual but not its label. In such cases the speaker must choose a descriptor which conveys those properties sufficient to identify the intended referent to the listener.
- (3) In both cases (1) and (2), descriptors were used on a specific reading. The third case assumes a non-specific reading for the descriptors. In this situation, the speaker intends to refer to the concept of something. Such concepts usually have no label available and a descriptor must be used.⁸

⁶ A descriptor is used *intensionally* to refer to either an individual concept or universal concept in the KB. Additionally, two intensional descriptors can be either definitionally equivalent or co-intensional depending on the equivalence between them.

⁷ Claiming that labels are unique is an oversimplification; for example John the Third, John Lennon, and John may all be proper names for the person John Lennon. One might consider them as isomorphic labels — abbreviations (if the context permits) for the full form.

⁸ Often the speaker is wanting to focus attention on a role rather than the filler of the role.

3.2. States of Affairs

To give a proper account of sentential constituents in opaque contexts, we must be able to refer to alternative, conjectured *states of affairs*.⁹ The need to conjecture about non-factual worlds arises for several reasons. A listener's knowledge of the current world (real or not) is incomplete, it can be uncertain, and does not fully determine future states. Thus, we are forced to make predictions about what is and isn't possible. As well we will often choose to discuss imaginary worlds and entities in them, with little relation to how the real world is. At any point in a discourse, some states of affairs are said to be "compatible", namely those which are plausible given the current state of affairs, the previous context, discourse situation, etc. Our notion of compatibility follows closely with that of Levesque (1983) in which "compatibility constraints" define *situations*, partial possible worlds in which only relevant information is asserted:

In particular, a situation may support the truth of some sentences and the falsity of others, but may fail to deal with other sentences at all. (Levesque 1983:199)

Compatibility can be formalized as a set of constraints which must hold for any state of affairs to be considered plausible. Of these constraints, some will be more applicable at a given time, some more susceptible to modification, and some more readily defeasible. A constraint then will have a particular *scope of effect* ranging from being fixed for the entire discourse, to applying only for the duration of a sentence, and a *priority* (or degree of relevance) in the current context. Because of these rankings, certain compatible states of affairs will be preferred.¹⁰ For example, in a fairytale setting one can imagine that the constraints reflecting some common physical laws would be implicitly suspended and so would have zero priority.

The listener considers a sentence with respect to a *current* state of affairs which is likewise determined by a set of constraints. A subset, which represent common facts and assumptions, are taken to be unconditionally *true* for the duration of the conversation. This current state of affairs is to designate what the agent explicitly believes the world *is* like given the present discourse situation. In contrast, *compatible* states of affairs refer to what the world *might* be like, although in most cases they will include the majority of the current world constraints as well. The current state is always compatible and is defined by those constraints having the highest priority and most immediate scope of effect.¹¹

Five sources of knowledge contribute to constraints on compatible states of affairs. These are ordered according to the permanency of the acquired constraint in the KB from always present to contingently present for only a subset of the dialogue.

(1) Linguistic Convention:

Within a particular dialect, the basic definitions and universals of the language are always taken to be unconditionally true and result in definitionally equivalent terms (*e.g.*, "bachelor" =*df* "unmarried man"). In our model these constraints are shared by all dialogue participants and so belong to permanent common knowledge.

(2) Non-linguistic Common Knowledge:

Certain facts (*e.g.*, Paris is the capital of France) are assumed to be known by everyone unless explicitly stated otherwise. This type of knowledge can be either domain-specific or of a general nature. Its scope of effect is minimally the duration of the conversation.

⁹ I prefer this terminology to the somewhat misleading term "possible worlds" of Kripke and others.

¹⁰ Such preferred states are useful for the understander to formulate his expectations for subsequent dialogue.

¹¹ Nothing is being asserted about which states are "real" and which are "fictitious". I contend that such real world correspondence should be embodied in the content of the constraints themselves.

- (3) **Mutual Assumptions of Speaker and Hearer:**
Some knowledge and assumptions shared by the speaker and listener are not established linguistically but are inferred from the previous discourse and the discourse situation. The exact set will depend on the particular domain and who the conversation participants are and will vary over the course of the dialogue.
- (4) **Explicit Definitional Knowledge:**
The speaker may establish the definitional equivalence of terms by explicitly defining them in the discourse. Alternatively, the speaker may stipulate certain conditions by the use of counterfactuals, or annotating the discourse (e.g., "This is a story...") which indicate how compatible states of affairs are similar to or different from the current state. Such constraints will initially have a high priority and their scope of effect will last until a change of topic.
- (5) **Explicit Linguistic Co-reference:**
The speaker may convey that terms are co-extensional by explicitly stating or pointing out that they have the same referent. These are updated more frequently than type (4) and assert nothing at the intensional level.

To understand opaque constructs, we must know the denotation of expressions not just for single states of affairs but for many possible states of affairs. Compatibility constraints are a mechanism to help us do this. In transparent contexts only co-extensionality and the current state of affairs are important. Modal and other opaque operators modify the compatibility constraints along particular dimensions of the states of affairs. These modifications must be recognized to properly evaluate constituents in an opaque scope.

3.3. Semantic Classification

States of affairs can vary along any of several dimensions relating to time, space, authenticity, necessity, speaker's conviction, holders of belief axioms, agents of descriptors, etc. These different dimensions, of *linguistic* significance, are comparable to similar types for which specialized logics have been developed. Thus we have *modal (alethic)* logic for reasoning about possibility and necessity, *deontic* logic for obligation and permission, *epistemic* logic for knowledge and belief, and *tense* logic for temporal reasoning.

Ideally, we could use these formal systems to model opacity in natural language. However, as was the case with the inferences discussed in section 2.2.1, the strict logical definitions often do not conform to our intuitions. Even though our operators for opacity may syntactically resemble operators of formal logics, their semantics will need to be modified to agree with our linguistic intuition.

The following categorizes the various opaque constructs according to their common semantic features. Some notation is introduced to represent the different semantics. A complete set will be given in subsequent chapters. Individuals relevant to the discourse are referred to as either:

- (a) the *speaker*, the originator of the sentence, or
- (b) the *listener*, the understander of the sentence, or
- (c) an *agent*, a third party referred to in the sentence.

The speaker and listener belong to the discourse situation while agents may or may not.

3.3.1. Agentless Sentential Opaque Contexts

Agentless sentential opaque constructs express a general opinion about an embedded sentential constituent. No agents are introduced; the crux of the problem is to interpret the

proposition relative to conjectured states of affairs.

Montague (1974) defined the underlying problem to be that propositions having the same truth value could not be intersubstituted, contrary to what *logical* intuition would predict. To account for this, Montague suggested that the truth value itself is not being referred to, but the *intension* of the sentence, defined to be a total function from possible world indices to extensions. Interpreted as such, only co-intensional constituents can be substituted in opaque contexts. Many computational linguists argue, and I concur, that the difficulties Montague envisions stem not from opacity but from using a truth-conditional analysis. By interpreting propositions as referring to the conditions they embody rather than to their truth value would appear to both overcome the substitution problem and be more linguistically appealing. Are Montague's suggestions then rendered useless? No, not quite, for he does provide some useful comments for these types of constructs, even for a non-truth-conditional scheme.

Agentless sentential opaque constructs assert statements *about* other statements. The conditions of standard (extensional) sentences are assumed to be satisfiable in the current state of affairs, whereas the satisfiability of propositions embedded in these opaque constructs concerns *hypothesized* states of affairs. Thus, there is a change in emphasis from knowing *if* a proposition *is* satisfied to knowing *what* conditions must hold for it *to be* satisfied (independently of whether those conditions are believed to currently be fulfilled or not). The stress is on the conditions and how they might come about rather than on their current satisfiability, which is more or less what Montague was getting at.

Montague suggests that states of affairs should be treated similarly to temporal affairs and for this uses "possible worlds" as the analog to points in time. This insight is also useful for NLP, especially since much work has been devoted to formulating temporal models and logics for AI. However, to apply these ideas to states of affairs we must first provide well-defined dimensions along which states of affairs can vary. As an initial step in this direction I subdivide this class of sentential opaque constructs into three groups (more are required) according to the type of dimension affected. For all of these, the NLP designer must decide how to represent conjectured states of affairs, and subsequently how to interpret the type of reference being asserted by a NP with respect to these states.

3.3.1.1. Objective Attitudinals

Objective attitudinals include not only expressions of logical possibility and necessity but also more general sentential modifiers which express *objective* attitudes (*e.g.*, degree of certainty or confidence) toward an embedded proposition. The speaker is ascribing the attitude not only to himself but to "everyone" in general. For example, "obviously" would be interpreted as "it is obvious to me and everyone else that . . .".

Syntactically this class includes phrases from:

Cleft Sentences: *It is possible that, It is necessary that, It must be the case that, . . .*

Modals: *may, could, might, . . .*

Sentential Adverbs: *apparently, obviously, necessarily, possibly, . . .*

Counterfactuals

Necessity and possibility are represented by the logical operators \Box and ∇ , respectively.

$\Box \phi$ expresses either that the proposition ϕ should be added as a permanent, non-defeasible constraint for all states of affairs, or that ϕ is easily provable from common knowledge, the explicit beliefs of the speaker, and current compatibility constraints.¹² Such a proof can rely

heavily on definitions of linguistic convention, as in

(3-16) *Necessarily* a bachelor is unmarried.

or on non-linguistic common knowledge, as evident in

(3-17) *Necessarily* the sun rises in the east.

$\nabla \phi$ corresponds to "It is possible that ϕ " and expresses that the proposition ϕ is *not inconsistent* with current compatibility constraints but does not actually assert ϕ . This corresponds to what Karttunen and Peters (1977) term *epistemically possible*, that is, possible in relation to what is currently known. Substitution of NPs is not restricted so long as the resulting phrase remains consistent with the KB.

Counterfactuals of the form "If ϕ , then ω " express an explicit constraint for the sentence. If the proposition ϕ is added to the current state, then the proposition ω should be easily provable. The addition of ϕ may render the current KB inconsistent; however this is not usually a problem, as ϕ will temporarily be given priority over other competing facts.

The more general types of attitudes such as "obviously" and "apparently" assert that a proposition either is part of common knowledge, or should be explicit¹³ (at least for the speaker) from the preceding discourse and the given situation. The speaker is often trying to make explicit some implicit beliefs of the listener. For example, "Apparently ϕ " asserts that the proposition ϕ cannot be proven for sure, but only a small amount of information is required to complete the proof.¹⁴ As a result the KB should assume ϕ is true but with less confidence than other facts.

These attitudes are represented by sentential operators. The major difficulty is defining the particular semantics for each construct in terms of knowledge base primitives. Many affect how the proof should proceed rather than what is proven. The reliance on conjectured states of affairs also increases the likelihood that a noun phrase will be used non-specifically.

3.3.1.2. Temporal Phrases

The semantic evaluation of temporal phrases depends on a time other than the time of utterance. The knowledge base needs to support some kind of temporal reasoning based on either partially ordered intervals or discrete points of time. This requirement is non-trivial in itself, especially since some tenses such as present perfect or progressive have no direct correlate in standard temporal logics (van Benthem 1985 : 10). For our purposes, the two operators F, for future time, and P, for past time, will suffice:

F ϕ = "it will be the case that ϕ "

P ϕ = "it has been the case that ϕ "

This class focuses on the time of reference. Embedded NP descriptors can have present or future realizability and can also be given a specific or non-specific reading. Likewise, at the sentential level we must determine the time when the specified conditions of a sentence are intended to apply. Consider:

¹² The logicians' use of "necessarily" to mean true in all *conceivable* states of affairs tends to be limited in NL and, if intended, sounds rather stilted and artificial. For this reason I relax the restriction that the embedded proposition must be true in all states of affairs to being true in a reasonable range of states of affairs implicitly agreed upon by the speaker and hearer, usually all compatible states of affairs.

¹³ Levesque (1983 : 198) makes the distinction between implicit and explicit beliefs. *Explicit* beliefs are those sentences actively held to be true by an agent whereas *implicit* beliefs describe what the world should be like given a set of explicit beliefs. A primary informational function of certain opaque contexts (e.g., "obviously") seems to be to make explicit some implicit beliefs of either the listener or another agent.

¹⁴ This requires having a measure of provability to know how "close" one is to a proof.

(3-18) Ronald Reagan is the President.

(3-19) The President *always* lives in the White House.

(3-20) Ronald Reagan *always* lives in the White House. (Mays 1985)

The third sentence (3-21) cannot be validly inferred from the preceding two. "Always" creates an opaque (temporal) context. While the descriptor "the President" is non-specific, "Ronald Reagan" is specific so the two cannot be substituted.

Similarly, the future tense can create an opaque context as in:

(3-21) A man in a hat *will* marry [a beautiful woman].

Within the scope of a future operator, the expected reading for an NP is non-specific and its descriptor is assumed to apply in the future (*e.g.*, some woman who will be beautiful in the future), whereas the transparent reading favours a specific interpretation with present application of the descriptor (*e.g.*, a particular woman who is now beautiful). A specific reading of the NP but with a future applicability of the descriptor must be explicitly stated by the speaker as in:

(3-22) John will marry a (certain) girl *who will be beautiful*.

3.3.1.3. Deontic Phrases

Deontic phrases express propositions of obligation and permission. The embedded proposition is usually considered relative to some future, hypothetical time.

must, should, ought to, may, . . .

Those deontic phrases expressing permission are to have the general form $ok(F \phi)$ to express that it is acceptable that the proposition ϕ be satisfied at some time in the future. Those expressing obligation will have the form $\neg ok(\neg F \phi)$ to express that it would be unacceptable if ϕ is not satisfied in the future. Both expressions are implicitly according to the speaker.

Deontic phrases are especially difficult to interpret with respect to a knowledge base; in contrast to modal constructs, the commitment is not to a "possible" state of affairs but rather to an "expected" one.

3.3.2. Sentential Opaque Contexts With Agents

This set of constructs introduces agents to whom propositions and/or descriptors are being attributed. To represent this, the sentential operator must be relativized to an agent as in $op(a, \phi)$ which attributes proposition $op(\phi)$ to the personal knowledge base of agent a (*i.e.*, $op(\phi) \in KB(a)$). For any NP descriptor, there can be ambiguity due to intensionality and, if the agent is not the speaker, due to descriptive content as well.

3.3.2.1. Propositional Attitudes

Propositional attitudes express an individual's attitude towards the content of a proposition.¹⁵ This attitude can be either that of the speaker himself, or of a second or third party to whom the speaker is attributing it. Propositional attitudes can be subclassified according to the sort of attitude expressed:

- (i) Epistemic:
certainly, I am certain that, . . .
supposedly, admittedly, definitely, . . .
believes, knows, thinks, . . . that

¹⁵ In contrast to the attitudes of part (3.3.1.1), these are *subjective*, that is, relative only to the speaker and the agent of the beliefs.

(3-23) *Supposedly* free trade will stimulate the manufacturing sector in Canada.

These statements report on the certainty that an individual has about the factual content of a proposition. In this way, they express an assertion about the content of an agent's (possibly the speaker's) personal knowledge base. Evidence for or against the proposition is restricted to facts believed to be contained in such a knowledge base¹⁶. These will be represented by a predicate of the form $pa(a, \phi)$, where pa is a propositional attitude predication, ϕ is a proposition, and a is an agent.

- (ii) Hypothetical Emotives:
wishes, wants, hopes that, . . .

(3-24) The police chief *expects* the incidence of drunk driving to decline.

These express attitudes towards the content of a proposition, but do not assert it to be true according to anyone's particular knowledge base. It is an attitude towards the future satisfiability of a proposition with the negation of the proposition presumed. These will have the general form of $ha(a, F \phi)$ where ha is a hypothetical attitude operator, a an agent, F the future operator, and ϕ some proposition.¹⁷

3.3.2.2. Reported Behaviour

A speaker can report on his own or another's behaviour on the basis of some empirical observation. Unless otherwise indicated, the proposition expressed is assumed by the speaker and the hearer to be factual.

(3-25) The boy heard that the blood supply was dangerously low.

- (i) Perceptual: *heard that, saw that, . . .*
- (ii) Verbal (i.e., indirect quotation): *said that, announced that, . . .*

The behaviour reported on is contained in a proposition; yet the speaker is not committed to the validity of its content. Freer substitution is permitted in these contexts, so long as the factual content relative to the speaker and hearer is preserved. Descriptors are generally attributed to the agent of the behaviour as well as the speaker. The general form is $rb(a, \phi)$ where rb is a predicate for reported behaviour, a is an agent, and ϕ is a proposition.

3.3.3. Intensional Transitive Verbs

Intensional transitive verbs take simple NPs or infinitive complements as objects. The primary concern is how to interpret descriptors that occur in such argument positions. A simple NP object of an intensional transitive verb often has a questionable "real-world" existence, is ambiguous with respect to intensionality, and can be attributed to various agents. In the following, a is an agent, δ is a descriptor whose referent is intensionally ambiguous, and \wedge marks the beginning of agent scopes.

- (i) Epistemic:
know, remember, conceive of, . . .

(3-26) John *knows* [the number of the safe].

(3-27) John *doesn't know* [the number of the safe].

¹⁶ Since an observer can never know for sure what facts another individual knows, we always assume an implicit prefix of either "The speaker believes that" or "The listener believes that".

¹⁷ Factive verbs such as *surprised, regrets, sorry to hear that, . . .* are also emotive but *do* presume that the embedded proposition is true and therefore are not considered as hypothetical emotives. Some factives such as *believes* fall under category (i), epistemic propositional attitudes, and thus are opaque.

Epistemic verbs express awareness by the subject of an entity described by the direct object NP. On a positive reading, (3-27) the "knowing" of the object NP refers to both an intension and extension; on a negative reading, (3-28) asserts that John doesn't know the value of the extension but presumably knows of the intension. These have the form $ep(a, \delta)$, where ep is an epistemic predicate.

- (ii) Emotive:
want, seek, need, expect, look for, . . .
 This type expresses a desire towards an entity. The object NPs are ambiguous with respect to intensionality, descriptive content, and real-world existence. The form $em(a, \delta)$, where em is an emotive predicate, is used to represent these cases.
- (iii) Stative:
becomes, is, resembles, . . .
 Stative verbs describe a mental or physical state or change of state. Most often these assert some type of equative relation between the subject and object descriptors. The form of these will be $eq(\delta_1, \delta_2)$, where eq is an stative predicate.
- (iv) Creative:
is writing, building, creating, . . .
 The creation of a specific object is described. The object NP is always given a specific reading but its existence is not entailed. We represent these with the form $cr(a, \delta)$, where cr is an creative predicate.
- (v) Performative:
ordered, commanded, . . .

(3-28) Harry *ordered* a new dishwasher.

When used transitively with a non-complement object NP, the NP is subject to ambiguity with respect to specificity. These take the form $pf(a, \delta)$, where pf is an performative predicate.

Many emotive verbs and performatives can also take infinitive complements. Any NP descriptor in the complement of an emotive verb can be opaquely ambiguous in the same way as simple NP objects, although the infinitive may help to resolve the ambiguity. If there is also another (simple) NP object (e.g., Mary *wants* [a man] to buy a coat), then that NP object is included in the intensional scope of the emotive verb. (Since it is actually the subject of the infinitive clause.) In contrast, when performatives take infinitive complements there must be an indirect object as well as the complement (e.g., Mary *ordered* [a man] to buy a coat.); this object is not considered in the intensional scope of the performative nor considered an agent for the infinitive clause.¹⁸

3.3.4. Direct Quotations and Titles

In sentences containing direct speech, titles (e.g., of books, people), etc. the speaker is referring to the form rather than the content of the phrase. Any type of constituent can be involved for direct quotations. Only substitutions which refer to the same quote are permitted. These substitutions usually resulted in a very strained sentence, as in the following example:

(3-29) John said '*Hi Sue*'.

(3-30) John said *the sentence consisting of the words 'Hi' and 'Sue', in that order.*

¹⁸ I do not think that performatives used with infinitive clauses create opaque contexts. In the sense that the complement may not yet have been carried out it is similar to a deontic phrase imposing an obligation for future action. However, it isn't clear that the clause and descriptions embedded in it need to be interpreted differently from the extensional case.

Similarly, for sentences with titles we can only substitute phrases referring to the same titles. These substitutions are also strained, as the following example illustrates:

(3-31) *Tiny Tim* was so-called because of his size.

(3-32) *The boy named Tiny Tim* was so-called because of his size.

For all sentences in this category, the quoted phrase may be attributed to an agent. In our notation, single surrounding quotes will demark referents of form as in ' ϕ ' where ϕ is some constituent.

3.3.5. Intensional Modifiers

This type of expression modifies the intension of a common noun and includes intensional adjectives, adverbs, and predicative verbs.

typical, good, superior, fake, . . .

(3-33) My mother served *mock* lobster for dinner.

The meaning of the modifiers depends on the meaning of what they are modifying. For this reason it is difficult to give a context independent semantics for them; I would hope that they could be categorized along similar lines as the states of affairs, but do not discuss them further.

3.4. Concluding Remarks

Table 2 provides a summary of the semantic classification proposed for opaque constructs and the general form of their associated representation. In the subsequent chapters we will stress in particular intensional transitive verbs, objective attitudinals, and propositional attitudes. Comments will be made about the other classes when appropriate.

SEMANTIC CLASSIFICATION		
Class	Example	General Form
I Agentless Sentential Opaque Constructs		
(i) Objective Attitudinals	<i>Obviously</i> IBM makes a lot of profit.	$\text{att } \phi$
(ii) Temporal Phrases	Sue <i>will</i> meet the captain of the football team.	$F \phi$
(iii) Deontic Phrases	You <i>may</i> call me Brenda.	$\text{ok}(F \phi)$
	Wendy <i>must</i> find a good job.	$\neg \text{ok}(\neg F \phi)$
II Sentential Opaque Constructs with Agents		
(i) Propositional Attitudes	Terry <i>believes that</i> your red car is brown.	$\text{pa}(a, \phi)$
(ii) Hypothetical Emotives	John <i>hopes that</i> her friend's dog wins the race.	$\text{ha}(a, \phi)$
III Reported Behaviour	John <i>heard that</i> a movie with Clint Eastwood is scary.	$\text{rb}(a, \phi)$
IV Intensional Transitive Verbs		
(i) Epistemic	David <i>doesn't remember</i> the location of the buried treasure.	$\text{ep}(a, \phi)$
(ii) Emotive	Ted <i>worships</i> a woman with golden hair.	$\text{em}(a, \phi)$
(iii) Stative	The computer scientist <i>is</i> a connoisseur of Chinese food.	$\text{eq}(\delta_1, \delta_2)$
(iv) Creative	There is a letter I <i>have to write</i> .	$\text{cr}(a, \phi)$
(v) Performatives	Lord Mountbatten <i>commanded</i> a huge battalion.	$\text{pf}(a, \phi)$
V Direct Quotation	" <i>Linguistics Made Easy</i> " is my favourite book.	ϕ

TABLE 2

CHAPTER 4

Descriptor Logic

In the preceding chapters, I have identified a set of linguistic constructs which exhibit common semantic irregularity and have suggested that they be processed in a uniform way. Traditionally, these constructs have variously been labelled as: *intensional*, *belief*, *attitude*, *opaque*, *modal* contexts¹. Indeed, many seem to display idiosyncratic semantic features which would necessitate a broad range of operators in a representation, thus destroying the apparent homogeneity of the class. I argue that this diversity among the constructs can be accounted for by evaluating *descriptors* according to intensionality, agents, time, or states of affairs. Introducing the descriptor notion preserves the homogeneity of the class while the dimensions provide enough detail to differentiate among the particular semantics of the constructs.

4.1. Montague Semantics Revisited

The linguistic philosopher Richard Montague, in his "PTQ" paper (1974), presents a fairly complete semantic system which provides an account of opaque constructs and their related phenomena. His system includes a set of corresponding syntax and semantic rules for translating the surface syntax into an intermediate representation based in intensional logic (IL). A further set of rules specifies a model-theoretic interpretation of this formula. We seek to exploit the intensional aspects of Montague semantics and incorporate them into a meaning representation (MR) which is compatible with our descriptor view of opacity.

Attracted by its formality and apparent generality, my initial intention was to construct a Prolog-compatible MR based on Montague's intensional logic scheme. In standard NLP systems all expressions are treated *extensionally*. This results in efficient and practical implementations, but they are not easy to adapt for opaque constructs. In contrast, Montague semantics provides a good formal treatment of opacity in which everything is given an *intensional* analysis. However, Montague's approach appeals to *logical* rather than *linguistic* intuition and does not represent the finer distinctions of opacity that I feel are necessary. To obtain standard (extensional) readings, meaning postulates are required, and the theoretical underpinnings make it difficult to modify for computational applications. The MR I have subsequently adopted tries to incorporate more linguistic intuition and common sense into the semantics of the various operators. As well, I stress the importance of descriptors as a key factor responsible for the problem of substitution in opaque contexts.

4.2. Descriptor Logic as a Meaning Representation

4.2.1. Basic Entities

Our meaning representation is constructed out of two fundamental entities: *terms*, which correspond to referable objects (*e.g.*, individuals, individual concepts, and universal concepts), and *predicates*, which assert relations among entities. Terms are either *constants*, *variables*, or *quantified terms*. Quantified terms introduce (individual) variables for subsequent predication and have the form "(Quantifier Variable : Restriction)". Predicates may be either *atomic*

¹ Collectively I group all of these under the one heading OPAQUE.

predicates (e.g., man), *modified predicates*, or *lambda abstracted predicates*. A modified predicate is of the form “{Modifiers, Predicate}”, where Modifiers is either an atomic predicate modifier or a list of such modifiers. Lambda abstraction of the general form “ λ Var (Formula)” allows us to construct (one-place) predicates out of an arbitrary combination of other predicates and terms. In general, predicates take a finite number of arguments which can be terms, other predicates, or formulas. Variables over predicates are also allowed.

Out of terms and predicates we build *formulas*. *Predicate formulas* result from the application of a predicate to term and/or other formula arguments. These can be *simple* or *conjoined* using the standard logical connectives (e.g., AND, OR). A *proposition* is a predicate formula which is preceded by zero or more quantifiers and has no unbound variables. A *restriction* is also a predicate formula but one introduced in a quantified term to constrain the range of the its associated variable; it may further include the separator “&&” (read as “such that”) to distinguish between the primary and the secondary information of the restriction.

Table 3 summarizes the type and notation of entities permitted in our MR. The notation for predicate application generally is $[Pred\ A0, A1, \dots, An]$, where *Pred* is the predicate and *A0*, ..., *An* its arguments.² If the result of the predicate application is a proposition it is given the notation $\langle Pred\ A0, A1, \dots, An \rangle$, where the arguments are as above.

4.2.2. General Form for NL Constituents

Our meaning representation encodes the literal content of a sentence as a proposition. It represents what was said and not necessarily what was intended. We choose to distinguish the following aspects of opaque contexts:

- (1) **Intensionality of Referents:**
NP descriptions are marked from the listener's perspective as either taking individuals, individual concepts, or universal concepts as referents according to whether definite descriptions are referential, attributive, or generic and whether indefinite descriptions are specific, non-specific, or generic.
- (2) **Opacity with respect to Descriptive Content:**
In sentences involving agents (usually as the subject of an opaque verb) we need to distinguish whether the descriptions of noun phrases are to be attributed to any of those agents or only to the speaker.
- (3) **Opacity with respect to Time:**
Certain opaque constructs imply that embedded descriptors can apply at some other time than the time of utterance (usually in the future).

The general form of a completed sentential clause will be a proposition of the form (Term-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.³ Usually its argument positions 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. Since quantified terms may have embedded sentential constituents in their Restriction (e.g., for relative clauses) and Predications may contain other sentential predicates, (Term-List) \langle Predication \rangle pairs can be nested inside of one another.

² This corresponds to $pred(A0, A1, \dots, An)$ of first-order predicate logics except we permit the predicate to be a variable.

³ The intention is that the Term-List provides the information to identify KB referents and the main predication asserts new information to be added to the KB.

MEANING REPRESENTATION

Entity	Example	Explanation
I Terms:		
Individual Variables	X, Y5, Z	Can be instantiated to individuals, or individual or universal concepts.
Individual Constants	he1, john, qe2	Extensional individuals.
Quantified Terms	(def X : [woman X])	Introduces individual variables.
II Predicates:		
Term Predicates	punker, racquet, confess	Relations taking $n \geq 0$ term or formula arguments.
Modified Predicates	{witty priest}	Intensifies the predicate.
Lambda-Abstracted Predicate	$\lambda X ([man X] \& [crazy X])$	Abstracts property from a complex formula.
Predicate Variables	P1, R2, Q	Instantiated to formula.
III Formulas:		
Predicate Formula	[seduce X,Y], [ugly X]	Result of predicate application.
Conjoined Formula	[sick X] OR [dead X]	Formulas conjoined by logical connectives.
Proposition	(def X : [man X]) <love mary,X>	Closed quantified formula.
Restriction	[man X] && (indef Y : [woman Y]) <love X,Y>	Property to restrict quantified variables.
IV Other:		
Predicate Modifiers	very, big	Modify other predicates.
Quantifiers	def, indef, class, label	Determiner of quantified term.
Restriction Separator	&&	Separator in restriction predicates read as <i>such that</i> .
Logical Connectives	AND, OR, COND	Conjoin formulas

TERMS + PREDICATES + OTHER = FORMULA

TABLE 3

Quantified terms arise either from quantified noun phrases in the input sentence or from stand-alone nouns (*e.g.*, mass nouns, plural count nouns). They have the general form “(Det X: R(X))” where Det is a quantifier corresponding to the determiner of the noun phrase, X is the variable introduced, and R(X) indicates restrictions on X. For purposes of illustrating the ideas of this thesis, we restrict ourselves to only three quantifiers: “indef”, “def”, and “class”⁴

⁴ In fact we require a larger set to include *the, some, many, ... etc.*, each marked as either definite or indefinite or class.

introduced by indefinite descriptions, definite descriptions, and class descriptions respectively. For mass and plural count nouns we use the *class* determiner in the quantifier position. We also allow a special *label* quantifier for use with proper nouns. Proper nouns are usually represented by constant terms and always have extensional referents. Labels in opaque contexts can only be ambiguous with respect to agents but not with respect to intensionality nor time of descriptor application. Thus, in opaque contexts which are ambiguous with respect to descriptive content, we may alternatively use a quantified term of the form “(label *c* : *Name*)”, where the *c* is an individual constant and *Name* is its label. Such quantified terms enable labels to be scoped relative to the agent scope marker.

Opacity with respect to agents and to time are treated as scope ambiguities while intensionality is marked as a binary distinction. In general, all quantified terms are left extraposed to the outermost term list. Those quantified terms marked as intensionally ambiguous may be prefixed by an intensional abstractor, *int_abs*, which indicates an intensional referent (*i.e.*, an individual or universal concept). Those quantified terms originating within the scope of the *agent scope marker*, $\hat{}$, may remain inside its scope and be evaluated relative to the possible agents at that point. Similarly, those quantified terms originating in the scope of the future operator, *F*, may stay inside its scope and thus implicate a future application of the descriptor.

4.2.2.1. Correspondence with Linguistic Constituents

Linguistic constituents are mapped into MR entities according to the following chart:

Sentence	→	Proposition
Proper Noun Phrase	→	(Labelled) Quantified Term
Pronoun	→	Individual Constant
Noun Phrase	→	Quantified Term
Common Noun	→	(1-place) Predicate
Adjective	→	Predicate Modifier
Determiner	→	Quantifier
Verb Phrase	→	Predicate Formula
Verb	→	Predicate
Auxiliary	→	Propositional Operator
Adverb	→	Predicate Modifier
Sentential Adverb	→	Propositional Operator
Prepositional Phrase	→	Predicate Modifier
Preposition	→	Predicate

4.2.3. Candidate Set Of Operators

To keep the inference rules and their application simple, we want to be as explicit as possible in the MR.⁵ This will require an enriched set of operators and predicates to capture the subtleties of opacity. (In all cases below ϕ is a proposition and *a* is an agent.)

4.2.3.1. Propositional Operators

Propositional operators are monadic or dyadic operators which map *propositions* into *propositions*. Term arguments embedded in the proposition can be either intensional or extensional.

⁵ This is an especially prudent strategy for logic programming, since the substitution rules can be reduced to a test for equality using unification.

- (1) **Modal Operators:** Modal operators affect the strength and scope of the propositional assertion.
 $\Box \phi$: ϕ must be satisfiable in all compatible states of affairs;
 $\nabla \phi$: ϕ must not be inconsistent with the current compatability constraints.
 Relevant Dimensions : Levels of Confidence; Degree of Belief
- (2) **Temporal Operators:** Temporal operators permit embedded descriptors to apply at times other than the time of utterance.
 $F \phi$: at some future time the conditions of ϕ will be satisfied;
 $P \phi$: at some time in the past the conditions of ϕ were satisfied.
 (also include fixed reference points like NOW, SINCE, UNTIL)
 Relevant Dimension: Time
- (3) **Conditional Operators:**
 $\omega \Rightarrow \phi$: if ω is added to the current state of affairs then ϕ should be satisfiable.
 Relevant Dimension: Explicit Conditions between States of Affairs

4.2.3.2. Term Predicates

- (1) **Epistemic Predicates:** These express relations between an agent and a proposition.
 $\langle \text{bel } a, \wedge \phi \rangle$: ϕ is believed by a to be satisfied in the current state of affairs;
 $\langle \text{nec } a, \wedge \phi \rangle$: a is committed to satisfying ϕ at some time in the future
 Dimension: Agent's belief axioms
- (2) **Believed-to-Exist Predicate:**
 $E(a, e)$: The entity represented by e is believed by agent a to have a referent in the current state of affairs.
- (3) **Common Knowledge Predicate:**
 $C(\phi)$: checks whether ϕ is in the common ground of the participants; *common ground* is the set of propositions that *any* participant is rationally justified in taking for granted by virtue of the previous discourse, the discourse situation, and mutual general knowledge (Karttunen and Peters 1977).
- (4) **Descriptors:** A *descriptor* is a predicate which provides a possibly incomplete specification of some referent, be it extensional or intensional.
 $[d1 X]$: indicates that the properties $d1$ 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.
 Dimension: Agent's description axioms

4.2.3.3. Other Things

- (1) **Agent Scope Marker:**
 \wedge : marks an *agent scope*. It can apply to a formula or term to indicate that any embedded descriptors must be evaluated with respect to the possible agents at the point where the scope of \wedge begins.
- (2) **Intensional Abstractor:**
 $\text{int_abs}(C, (\text{Quantifier } Var : \text{Description}))$: asserts that the quantified term, in the second argument position, is to have an intensional referent, returned in C . If C is subsequently predicated, then its referent is a universal concept; if the Var instead is predicated then the referent is an individual concept. Without the int_abs , the use of the Var refers to an extensional individual.

4.3. Examples

Before discussing the “finer” points of our MR, we should illustrate our notation in action. The following example provides the representations of its four readings plus an extensional counterpart.

(4-1) The man *wants* [a coat like Bill's].

- (1) Extensional referent, opaque with respect to descriptive content:

There is a coat the man wants and he describes it as being like Bill's coat.

(def Y : [man Y]) <want Y, ^ (indef X : [coat-like-bills X]) >

- (2) Extensional referent, transparent with respect to descriptive content:

There is a coat the man wants and the speaker describes it as being like Bill's coat.

(def Y : [man Y]) (indef X : [coat-like-bills X]) <want Y, ^X>

- (3) Intensional referent, opaque with respect to descriptive content:

The man wants something he describes as being a coat like Bill's.

(def Y : [man Y]) <want Y, ^ int_abs(C, (indef X : [coat-like-bills X]))>

- (4) Intensional referent, transparent with respect to descriptive content:

The man wants something the speaker would describe as a coat like Bill's.

(def Y : [man Y]) int_abs(C, (indef X : [coat-like-bills X])) <wants Y, ^X>

- (5) Non-Opaque Sentence:

(4-2) The man *buys* a coat like Bill's.

(def Y : [man Y]) (indef X : [coat-like-bills X]) <buy Y, X>

4.4. Discussion of Decisions

4.4.1. Opaque Scopes

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

4.4.2. Determiner Translation

In computational linguistics, NL determiners are commonly represented by three-place quantifiers in the MR of the general 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 NL determiners rarely have a direct correlation with the existential and universal quantifiers of first-order logic. In many of the MRs used with logic grammars⁷, determiners provide the basic structure of the MR formula. The determiners are translated into quantifiers and are all left-extraposed (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.

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

⁷ Dahl (1981) for example.

Another approach focuses more on the main verbs by first translating them into predicates and subsequently finding appropriate fillers for their arguments which 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 " $\lambda P \lambda Q (\exists x P\{x\} \text{ AND } Q\{x\})$ " which translates the determiner "a". Only quantified NPs in the subject position of a sentential clause will have the main predication filled. Those in the object/complement positions leave the Q predication unbound.

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 temporally or agent opaque) are left-extraposed and assimilated into a single list structure followed by a single main predication.

The different versions for determiner translation can be compared using the sentence:

(4-3) Each man saw a woman.

I use a similar notation based on that of Dahl (1981) to translate each (with the original author's syntax in square brackets).

(1) Main predication *always* in rightmost predicate position:

$\text{each}(X, \text{man}(X), \text{exists}(Y, \text{woman}(Y), \text{saw}(X, Y)))$ (Dahl 1981)

(2) Main predication in subject restriction, complement predicate unbound:

$\text{each}(X, \text{man}(X) \rightarrow \text{exists}(Y, \text{woman}(Y) \text{ AND } \text{saw}(X, \lambda Q [Q\{Y\}]))$
 $[\forall x \text{ man}(x) \rightarrow \exists y [\text{woman}(y) \text{ AND } \text{saw}(x, \lambda Q [Q\{y\}])]]$ (Montague 1974)

(3) Main predication separate:

$\text{each}(X, \text{man}(X)) \text{ exists}(Y, \text{woman}(Y)) : \text{saw}(X, Y)$
 $[(\text{def } X : [\text{man } X]) (\text{indef } Y : [\text{woman } Y]) <\text{saw } X, Y>]$ (Fawcett 1985)

4.4.3. Adjectives and Adverbs

Adjectives and adverbs are represented as predicate modifiers. Curly brackets enclosing a list of predicates signify modified predicates as in:

$\{\text{passionately}, [\text{loves}, \text{mary1}]\}$.

The first argument supplies the modifiers which can be either an atom or a list. This notation allows the order of the modifier application (especially adjectives) to be left ambiguous. For example, "bright, red ball" in the regular predicate notion would be expressed as one of:

- (i) $[[\text{bright}, \text{red}], \text{ball}]$, or
- (ii) $[\text{bright}, [\text{red}, \text{ball}]]$, or
- (iii) $[[\text{bright AND red}], \text{ball}]$.

To avoid selecting one, I just use the notation " $\{[\text{bright}, \text{red}], \text{ball}\}$ " which then is open to all the above readings.

4.4.4. Iterated Opaque Contexts

Our notation should remain useful in iterated contexts to handle multiple agents and ensure the proper inheritance of opacity or transparency from constituents to parents. The following example indicates how they appear in our notation.

(4-4) Pat believes that [a man] wants to meet [Pete's wife].

(1) Standard Case:

Pat believes that a particular man wants to meet a certain woman who is Pete's wife.

(indef X: [man X]) (def Y: [petes-wife Y])
 <bel pat, ^ <want-event X, ^ <meet X, Y>>>

- (2) Intensional on Both:

Pat believes that some man wants to meet whoever is Pete's wife.

int_abs(C1, (indef X: [man X])) int_abs(C2, (def Y: [petes-wife Y]))
 <bel pat, ^ <want-event X, ^ <meet X, Y>>>

- (3) Intensional on One but not Both:

- (a) *Pat believes that some man wants to meet the particular person who is Pete's wife.*

int_abs(C1, (indef X: [man X])) (def Y: [petes-wife Y])
 <bel pat, ^ <want-event X, ^ <meet X, Y>>>

- (b) *Pat believes that a particular man wants to meet whoever is Pete's wife.*

(indef X: [man X]) int_abs(C2, (def Y: [petes-wife Y]))
 <bel pat, ^ <want-event X, ^ <meet X, Y>>>

For each of these we must also consider opacity with respect to descriptive content. The scoping in the following examples does not commit us to a choice of agent but represents the maximal set of possibilities. That is, "a man" can be attributed to Pat and the speaker while "Pete's wife" can be attributed to the man, Pat and the speaker:

- (1)' Standard Case:

Pat believes that a particular man wants to meet a certain person who Pat and that man would describe as "Pete's wife".

<bel pat, ^ (indef X : [man X]) <want-event X, ^ (def Y : [petes-wife Y])
 <meet X, Y>>>

- (2)' Non-Specific on Both:

Pat believes that some man wants to meet whoever she would describe as "Pete's wife".

<bel pat, ^ int_abs(C1, (indef X : [man X])) <want-event X,
 ^ int_abs(C2, (def Y : [petes-wife Y])) <meet X, Y>>>

- (3)' Non-Specific on One:

- (a) *Pat believes that some man wants to meet the particular person who both Pat and the unknown man would describe as "Pete's wife".*

<bel pat, ^ int_abs(C1, (indef X : [man X])) <want-event X,
 ^ (def Y : [petes-wife Y]) <meet X, Y>>>

- (b) *Pat believes that a particular man wants to meet whoever Pat and that man would describe as Pete's wife.*

<bel pat, ^ (indef X : [man X]) <want-event X,
 ^ int_abs(C2, (def Y : [petes-wife Y])) <meet X, Y>>>

Treating ambiguity with respect to descriptive content as a scope phenomenon ensures the correct inheritance from children to parent constituents, and does not conflict with the representation of the intensionality of the referent.

4.4.5. Non-Opaque Scoping

If ascribing descriptors to agents is handled by a scoping mechanism, we must define how these scopes should interact with the scoping of standard determiners and prepositional phrases.

For example, the sentence

(4-5) Every man in our class *wants* a beautiful woman.

can have three potential readings:

For each man there is some particular beautiful woman that he wants.

(all X : [man-in-our-class X]) (indef Y : [beautiful-woman Y])
 $\langle \text{want-entity } X, \hat{Y} \rangle$

There is a certain beautiful woman whom every man wants.

(indef Y : [beautiful-woman Y]) (all X : [man-in-our-class X])
 $\langle \text{want-entity } X, \hat{Y} \rangle$

Each man wants (his concept of) a beautiful woman.

(all X : [man-in-our-class X])
 $\langle \text{want-entity } X, \hat{\text{int_abs}}(C, \text{indef } Y : [\text{beautiful-woman } Y]) \rangle$

The first two readings stem not from opacity but from the relative scoping of standard quantifiers. The third, intensional reading permits only one translation. This agrees with our intuition since it would be rather difficult to have each man having the same concept of a beautiful woman.

4.5. Descriptor Relations

When discussing and comparing descriptors, it is useful to define some notation for the type of identity among descriptors. These relations are introduced as convenient "syntactic sugar" which abbreviate the longer forms using complete predicates with agents and intensional abstractors.

(i) $P =_d Q$:

Descriptors P and Q are definitional equivalents, as in "stockowner" $=_d$ "shareholder". This should be part of (or added to) common knowledge (i.e., independent of any agents) and holds for all states of affairs.

$P =_d Q$ is defined to be equivalent to:

$\text{int_abs}(C1, (\text{Det1 } X1 : [P X1])) \text{int_abs}(C2, (\text{Det2 } X2 : [Q X2])) \square \langle C1 = C2 \rangle$

which should be read as:

The concept described by property P is always the same concept as that described by property Q.

(ii) $P =^{\wedge} Q$:

Descriptors P and Q are co-intensional; they co-refer in all contexts. $P =^{\wedge} Q$ is an abbreviation for:

$\text{int_abs}(C1, (\text{Det1 } X1 : [P X1])) \text{int_abs}(C2, (\text{Det2 } X2 : [Q X2])) \square \langle X1 = X2 \rangle$

to be read as:

The concept identified by property P will always apply to the same objects as the concept identified by property Q.

(iii) $P =_x Q$:

The fact that expressions P and Q are co-extensional in the current state of affairs is part of the speaker and listener's mutual knowledge.

$P =_x Q$ is defined as:

$(\text{Det1 } X1 : [P X1]) (\text{Det2 } X2 : [Q X2]) \langle X1 = X2 \rangle$

which can be paraphrased as:

In the current state of affairs, the object identified by property P is the same object identified by property Q.

Using the relation for co-extensionality, the sentence

(4-6) The Morning Star is the Evening Star.

might be expressed as

(4-7) [Morning Star] =_x [Evening Star]

whereas the sentence

(4-8) *Necessarily* the Morning Star is the Evening Star.

would be translated as

(4-9) [Morning Star] =[^] [Evening Star]

All three types of descriptor relations should ideally be considered with respect to an agent. Agents seem most important for substitution of co-extensional and co-intensional descriptors, so we provide alternative notation for these relative to agents. Knowledge of definitional equivalence we assume to be global. Thus we define:

(iv) $P =_a Q$:

This relation asserts that the agent, a , would use either of the descriptors P or Q extensionally to describe the same entity.

(v) $\Box (P =_a Q)$:

This relation asserts that in all contexts the agent, a , would use either of the descriptors P or Q to describe the same entity.

Our multiple notation for equivalence relations is:

Relation	Agentless	With Agents
Co-extensions	$X =_x Y$	$X =_a Y$
Co-intensions	$X =^ Y$	$\Box (X =_a Y)$
Definitional Equivalents		$X =_{df} Y$

4.5.1. Merging of Descriptors

If two descriptors independently describe an extensional object, then both together certainly will identify the same object. This can be expressed as an axiom in our notation:

$$\begin{aligned} & \forall P (\text{Det } X : d1) (\text{Det } Y : d2) \langle X = Y \rangle \text{ AND } \langle P X \rangle \\ & \rightarrow (\text{Det } X : d1 \text{ AND } d2) \langle P X \rangle \end{aligned}$$

This corresponds to extensional node merging of Maida and Shapiro (1982) and the extensionalizing operation of Nishida (1983). Such a merge also can be done for co-intensional descriptors, but this is a more difficult operation for the knowledge base. Expressed in our notation this is:

$$\begin{aligned} & \forall P \text{ int_abs}(C1, (\text{Det } X : d1)) \text{ int_abs}(C2 (\text{Det } Y : d2)) \Box \langle X = Y \rangle \text{ AND } \langle P X \rangle \\ & \rightarrow \text{int_abs}(C, (\text{Det } X : d1 \text{ AND } d2)) \langle P X \rangle \end{aligned}$$

Merging becomes expensive to implement if the referents are already in the KB. A lot of other assertional and definitional information may be attached to them which must be "sorted out" before the nodes can be merged.

CHAPTER 5

Substitution of Equals

5.1. Range of Substitution

The failure of substitution of equivalent phrases appears to be a gradable notion; the degree of substitution allowed varies with the type of opaque construct under consideration. We can think of a scale of substitutivity, with the lower bound being a strictly *de dicto* reading in which only substitutions of phrases referring to the form of the original phrase are permitted and the upper bound a strictly *de re* reading in which co-extensional phrases can be substituted in any context.

As we saw in chapter 3 (section 3.3.4), sentences containing direct quotations or titles provide examples where only very limited substitution can take place. The speaker is reporting on the form as well as the content of the quoted phrase. Only phrases referring to the identical form can be interchanged.

For sentences of indirect discourse, the speaker is not bound by the exact words of the reported proposition. Relative to the mutual knowledge of the speaker and hearer, the speaker is free to substitute any phrases which preserve the factual accuracy of the original utterance. Suppose John says "Bill is a bachelor", then it is acceptable for a speaker to subsequently state that

(5-1) John *said that* Bill is [unmarried].

since only information conveyed in the original utterance is being paraphrased.

Even with sentences of propositional attitude, certain descriptors can be substituted for, provided the content of the proposition being ascribed (again relative to the speaker and the hearer) is not affected. However, this is not the only criterion which determines terms which are interchangeable.

(5-2) Bill *thinks that* [Queen Elizabeth] is a lovely lady.

(5-3) Bill *thinks that* [the Queen of England] is a lovely lady.

(5-4) Bill *thinks that* [the titular head of the Church of England] is a lovely lady.

In the above example the speaker and hearer assume implicitly that, since the filler of the role "Queen of England" is not likely to change within the time of their conversation and that they are both aware of who fills that role, it is acceptable to allow substitution of the role with the filler and vice versa. Thus, sentence (5-3) can be inferred from (5-2). However, suppose that Queen Elizabeth is also the titular head of the Church of England. To substitute this phrase instead of "the Queen of England" as in (5-4) seems to attribute further knowledge to Bill than was in the original statement.¹ The criterion then appears to be that the descriptors not only be

¹ A similar discrepancy between role and role-filler arises in non-opaque contexts. Hofstadter (1982) presents a school scenario in which the school secretary instructs a student to "Have Mr. Talbot sign your form". Mr. Talbot is the principal of the school. Both realize that the secretary is in fact referring to Mr. Talbot in his role as principal of the school. If someone else's name were substituted, say Mr. Smith the janitor, in place of Mr. Talbot, the student would probably look very puzzled and require further explanation. It is not clear why the janitor's signature should be required. In the initial case, the role filler is being used to refer to the role itself. This fails in the second

coreferential, but coreferential according to the agent of the propositional attitude.

In contrast to characterizations of classical logic, the problem of substitution in opaque contexts stems from the failure to recognize how descriptors relate, and **not** 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.

5.2. The Substitution Phenomenon

We now examine more closely the criteria for substitution in opaque contexts. The clause containing the noun phrase to be substituted for is called the *target* clause and the noun phrase itself the *target* descriptor. The knowledge-base statements that authorize such substitution will be termed *background assumptions* (BASS).²

Theories of formal logic dictate that substitution should be permitted, but we must question whether there are *linguistic* reasons for it. At least two seem important for NLP systems. In both the NLP system plays the role of the listener, trying to emulate his understanding of and subsequent reasoning about an input target sentence.

- (i) **Extracting Maximum Information:** The listener wants to get the maximum information out of the target and assimilate it with his previous knowledge in a consistent and economical way. Both aspects, extracting the new information, and assimilating it into the KB, require understanding how descriptors in the target "relate" to those in the BASS. It may be more straightforward to reason with a "known" descriptor from the BASS substituted in the target than with the "new" descriptor.
- (ii) **Re-Phrasing Information for a Third Party:** After hearing and understanding the target the listener may decide to convey its content to a third party. The listener will want to preserve the factual accuracy but will have to use descriptors that the *third party* will understand. These often will not be the same as those in the target sentence. Thus the (original) listener will need to translate descriptors of one agent into descriptors that another agent (the third party) is aware of; an ideal application for substitution rules.

5.2.1. A Detailed Example

Often in texts and papers discussing opacity, three sentences are given without stating explicitly their status as to who knows them, spoke them, or is trying to understand them. It is simply put that from the first two sentences (read opaquely) it is not valid to infer the third.

To better understand the underlying assumptions and the role of descriptors we now consider an extended example:

(5-5) [The Dean] is [the Chairman of Imperial Oil].

(5-6) John is *looking for* [the Dean].

(5-7) John is *looking for* [the Chairman of Imperial Oil].

The first sentence (5-5) belongs to the BASS of the listener. He may believe about (5-5) that

- (a) both he and the speaker are aware of the statement expressed, that is, it belongs to mutual knowledge³, or

case since the referent does not fill the expected role.

² These usually assert various kinds of descriptor equivalence.

³ This could be established earlier on in the discourse text or inferred from it.

- (4) Two terms may be made equivalent explicitly in the preceding discourse. For example, the statement "Brian Mulroney is the Prime Minister" will authorize substitution in most subsequent discourse such as from

(5-8) [Brian Mulroney] is my hero.

to

(5-9) [The Prime Minister] is my hero.

However, a statement like "Clark Kent is Superman" only permits substitution of the two phrases in subsequent *extensional* contexts, since from the sentence

(5-10) [Superman] is my hero.

we cannot substitute to obtain

(5-11) [Clark Kent] is my hero.

We must be careful to note what level of authorization is being given and permit substitution accordingly.

5.3.1. Copular Relations and their Descriptions

As part (4) above indicates, equivalence among descriptors is often asserted linguistically by using copula verbs. It is therefore important to interpret them correctly. Mallery (1985) provides an insightful scheme for classifying noun phrases in the subject and object positions of copula verbs. In his semantic model (*i.e.*, RELATUS) quantified noun phrases inherit the definiteness of their determiner⁴ which belongs to one of three states: *indefinite* (*e.g.*, "a"), *definite* (*e.g.*, "the"), or *class* (*e.g.*, "all").⁵

These *indefinite*, *definite*, and *class* noun phrases are then interpreted in context as referring to INDIVIDUALs or UNIVERSALs. UNIVERSALs ultimately correspond to class or frame objects while INDIVIDUALs correspond to instances of such objects. The interpretation of NPs as UNIVERSALs or INDIVIDUALs is governed by the major verbs of a clause and the grammatical relations among clauses. For example, when two indefinite NPs are connected by a copula, both are interpreted as UNIVERSALs. It turns out that copula verbs (in particular "be") assert one of four types of links that are useful for constructing a class structure in the knowledge base:

Ordinary Classification	(INDIVIDUAL ↔ UNIVERSAL)
Generic Classification	(UNIVERSAL ↔ UNIVERSAL)
Identity Relation	(INDIVIDUAL ↔ INDIVIDUAL)
Quality	(INDIVIDUAL/UNIVERSAL ↔ ADJECTIVE)

If more than one type of link is possible for a verb, it is necessary to look at previous referents in the discourse focus to resolve the ambiguity.

Mallery's scheme is fairly compatible with my work, although I believe he omits some finer points which are of significance for opaque contexts. He does not differentiate between specific and non-specific readings of NP descriptions; doing so adds some extra complexities (and ambiguities) to his table of NP quantification (Mallery 1985:38). Instead of three states of definiteness, there would be five: *definite-referential*, *definite-attributive*, *indefinite-specific*,

⁴ The definiteness of determinerless NPs is established using other parameters such as number or type of noun.

⁵ Mallery does not include partitive determiners such as "most", "few" which refer to subsets of classes. It is not obvious how these could be incorporated into his scheme.

indefinite-non-specific, and *class*.

As well, I would argue for two further types of classificational links for copula verbs. This stems partially from the previous suggestion, but also from the ways in which descriptors can be used. As we have previously noted, two NP descriptors may:

- (1) co-refer in the current context (Co-Extensions),
- (2) co-refer in *all* contexts (Co-Intensions),
- (3) express the *same concept* (Definitional Equivalents).

A new relation is needed to assert equivalence of type (3); for this I suggest a new copular link between UNIVERSALS which corresponds to my *definitional equivalence*.

The second link I suggest distinguishes between one-way and two-way generic classification (analogous to " \rightarrow " and " \leftrightarrow " of standard logic). Two-way generic classification, which I rename *intensional equivalence*, handles cases of type (2) above where the interchange of subject and object does not affect the relation being asserted. One-way generic classification handles cases where the subject is a strict subclass of the object class; interchanging subject and object inverts the asserted relation.

This modified framework for classifying NP descriptions in argument positions of copula verbs is summarized in the next two pages. (CLASS type determiners are not included.) It will be later incorporated into our NLP semantic component.

COPULAR RELATIONS

I. Ordinary Classification ($i \in C$)

INDIVIDUAL R UNIVERSAL

DEF-Ref	INDEF-Nsp	The Prime Minister is a Conservative.
INDEF-Sp	INDEF-Nsp	A neighbour's dog is an airedale.

II. Generic Classification ($C1 \subseteq C2$)

UNIVERSAL R UNIVERSAL

INDEF-Nsp	INDEF-Nsp	A dog is a mammal.
DEF-Nref	INDEF-Nsp	The Prime Minister is an elected official.

III. Extensional Equivalence (Identity) ($i = j$)

INDIVIDUAL R INDIVIDUAL

DEF-Ref	DEF-Ref	The Prime Minister is the President of Alcoholics Anonymous.
INDEF-Sp	DEF-Ref	A man I met yesterday is the President of Alcoholics Anonymous.

IV. Intensional Equivalence $\forall x (C1(x) \equiv C2(x))$

UNIVERSAL R UNIVERSAL

DEF-Nref	DEF-Nref	The Queen of England is (always) the Commander of the Navy.
INDEF-Nsp	DEF-Nref	A puma is the fastest mammal.

V. Definitional (Conceptual) Equivalence ($C1 \equiv C2$)

UNIVERSAL R UNIVERSAL

INDEF-Nsp	INDEF-Nsp	A bachelor is an unmarried man.
INDEF-Nsp	DEF-Nref	A coxswain is the person who calls the strokes on a rowing scull.

VI Extensional Predication $\text{Pred}(\text{Arg})$

INDIVIDUAL R ADJECTIVE

DEF-Ref	ADJ	The Prime Minister is foolish.
INDEF-Sp	ADJ	A neighbour's dog is fierce.

VII. Intensional Predication $\text{Pred}(\text{Arg})$

UNIVERSAL R ADJECTIVE

DEF-Nref	ADJ	The Prime Minister is well-paid.
INDEF-Nsp	ADJ	A lion is fierce.

NOTATION FOR COPULAR RELATIONS

I. ORDINARY CLASSIFICATION:

(def X: [prime_minister X]) <conservative X>

(indef X: [dog_of_neighbour X]) <airedale X>

II. GENERIC CLASSIFICATION:

int_abs(C, indef X: [dog X]) <mammal C>

int_abs(C, def X: [prime_minister X]) <elected_official C>

III. EXTENSIONAL EQUIVALENCE:

(def X: [prime_minister X]) (def Y: [president_AA Y]) <X = Y>

IV. INTENSIONAL EQUIVALENCE:

int_abs(C1, def X: [queen_of_Eng X]) int_abs(C2, def Y: [commander_navy Y]) \square <X = Y>

int_abs(C1, indef X: [head_competition X]) int_abs(C2, def Y: [best_racquet Y]) \square <X = Y>

V. DEFINITIONAL EQUIVALENCE:

int_abs(C1, indef X: [bachelor X]) int_abs(C2, def Y: [unmarried_man Y]) <C1 = C2>

VI. EXTENSIONAL PREDICATION:

(def X: [prime_minister X]) <foolish X>

VII. INTENSIONAL PREDICATION:

int_abs(C, def X: [prime_minister X]) <well_paid X>

int_abs(C, indef X: [lion X]) <fierce C>

Table 5.

5.4. Minimal Requirements for Substitution

If NP descriptors are to be substituted in opaque contexts, 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. For transparent contexts, a descriptor always has an extensional referent, refers at the time of utterance, and is ascribed to the speaker only. But, in opaque contexts, any or all of these factors may vary. We must establish the "level" of each factor in the target sentence and then determine if axioms in the BASS authorize substitutions at that level. That is, we must relate the intensionality, time, and agent of the *descriptor equivalence* asserted in the BASS to the intensionality, time, and agent of the target *descriptor*, and then assert the intensionality, time, and agent of the descriptors in the resulting clause (after any substitutions). We will look at each aspect in turn and then consider their interactions.

A typical substitution replaces the target descriptor, *d1*, with an equivalent descriptor, *d2*, from the BASS but otherwise preserves the form of the target sentence (*i.e.*, RESULT = TARGET [*d1/d2*]). We can also get atypical substitution, in which the modality of the BASS affects the modality of the target relation. The assertion in the resulting clause may have a different modality as well as substituted descriptors (*i.e.*, RESULT = NEW_TARGET [*d1/d2*]). These inferences rely on linguistic intuition for their validity rather than on deductive principles of logic. Thus, their status remains dubious and the correctness of the inference is often debatable. Nevertheless we try to provide some account of them.

5.4.1. Intensionality of the Descriptor

In many opaque constructs, no agents are present and yet substitutions predicted by logical inference rules still fail. A common reason is that the intensionality of the descriptor equivalence asserted in the BASS is not compatible with the intensionality of the target descriptor.

The typical rules for substitution according to the intensionality of the descriptor are:

- (1) If the target descriptor is extensional, then co-extensional, co-intensional, or definitionally equivalent descriptors in the BASS can be substituted.
- (2) If the target descriptor is intensional (individual concept), then co-intensional or definitionally equivalent descriptors in the BASS can be substituted.
- (3) If the target descriptor is definitional (universal concept), then definitionally equivalent descriptors in the BASS can be substituted.

In all of these cases, the result will be the unchanged target assertion with the BASS descriptor substituted for the target descriptor at the level of intensionality of the target descriptor.

The rule for atypical substitution is:

- (4) If the modality of the BASS equivalence is strictly "weaker" than the modality of the target assertion, substitution can still be carried out according to rules (1) to (3) but the result will have the weaker modality applied to the original target assertion.

Consider comprehensive example A:

Background Assumptions:

I	<i>The Morning Star is the Evening Star.</i>	(def X: [morning_star X]) (def Y: [evening_star Y]) <X = Y>
II	<i>Possibly the Morning Star is the Evening Star.</i>	(def X: [morning_star X]) (def Y: [evening_star Y]) ∇ <X = Y>
III	<i>Necessarily the Morning Star is the Evening Star.</i>	(i) int_abs(C1, def X: [morning_star X]) int_abs(C2, def Y: [evening_star Y]) \square <X = Y> (ii) int_abs(C1, def X: [morning_star X]) int_abs(C2, def Y: [evening_star Y]) <C1 = C2>

Target Sentences:

A	<i>The Morning Star rises.</i>	(def X: [morning_star X]) <rises X>
B	<i>Possibly the Morning Star rises.</i>	(def X: [morning_star X]) ∇ <rises X>
C	<i>Necessarily the Morning Star rises.</i>	(i) int_abs(C, def X: [morning_star X]) \square <rises X> (ii) int_abs(C, def X: [morning_star X]) \square <rises C>

Substitutional Inferences:

A + I	<i>The Evening Star rises. (Rule 1)</i>	(def Y: [evening_star Y]) <rises Y>
A + II(i)	<i>?Possibly the Evening Star rises. (Rule 4)</i>	(def Y: [evening_star Y]) ∇ <rises Y>
A + III(i)	<i>The Evening Star rises. (Rule 1)</i>	(def Y: [evening_star Y]) <rises Y>
A + III(ii)	<i>The Evening Star rises. (Rule 1)</i>	(def Y: [evening_star Y]) <rises Y>
B + I	<i>Possibly the Evening Star rises. (Rule 1)</i>	(def Y: [evening_star Y]) ∇ <rises Y>
*B + II	<i>*Possibly the Evening Star rises. (Rule 4)</i>	(def Y: [evening_star Y]) ∇ <rises Y>
B + III(i)	<i>Possibly the Evening Star rises. (Rule 1)</i>	(def Y: [evening_star Y]) ∇ <rises Y>
B + III(ii)	<i>Possibly the Evening Star rises. (Rule 1)</i>	(def Y: [evening_star Y]) ∇ <rises Y>
C(i) + I	<i>nil</i>	nil
C(i) + II	<i>nil</i>	nil
C(i) + III(i)	<i>Necessarily the Evening Star rises. (Rule 2)</i>	int_abs(C, def Y: [evening_star Y]) \square <rises Y>
C(i) + III(ii)	<i>Necessarily the Evening Star rises. (Rule 2)</i>	int_abs(C, def Y: [evening_star Y]) \square <rises C>
C(ii) + I	<i>nil</i>	nil
C(ii) + II	<i>nil</i>	nil
C(ii) + III(i)	<i>nil</i>	nil
C(ii) + III(ii)	<i>Necessarily the Evening Star rises. (Rule 3)</i>	int_abs(C, def Y: [evening_star Y]) \square <rises C>

By pairing each of the target sentences, A, B, and C, with each of the background assumptions, I, II, and III, and then applying the above rules, we can generate a set of acceptable substitutional inferences. (The pairs of target and background assumptions which do not authorize substitution are marked with "nil" in the examples.)

Background assumption I asserts co-extensionality of the descriptors "Morning Star" and "Evening Star", assumption III(i) asserts their co-intensionality, and assumption III(ii) asserts their definitional equivalence.⁶ Assumption II asserts a weaker version of descriptor equivalence, it is only *possible* that the two descriptors are co-extensional (in some state of affairs). Similarly, for the target translations, the descriptor "Morning Star" in A and B is used extensionally, while in C it is intensional; C(i) refers to an individual concept, and C(ii) to a universal concept.

According to rule (1), background assumptions I, III(i), and III(ii) each authorize substitution in target clauses using either of the descriptors "Morning Star" or "Evening Star" extensionally. Thus, substitution of the descriptor "Evening Star" is permitted in A and in B at the extensional level.

Rule (2) states requirements for substituting individual concepts. Thus, background assumptions III(i) and III(ii) each authorize co-intensional substitution in the target clause C(i). Likewise, rule (3) permits us to use background assumption III(ii) to substitute definitional equivalents in the target clause C(ii).

Since the equivalence in background assumption II is only *possible*, it authorizes only atypical substitution at the extensional level. If we follow rule (4), we can substitute descriptors in target A and, potentially, in target B. However, the modality of the resulting target clause will be "weakened". The inference to be performed in A can be paraphrased as:

*Since the Morning Star may be the Evening Star (II), and
since the Morning Star rises (A), then
it is possible that the Evening Star rises (A+II).*

The equivalence assertion of assumption II is modally weaker than the assertion of target A, so we can substitute. *Logical* justification for such an inference is questionable, but it does seem intuitively correct.⁷

A similar substitution in target B is more dubious. In addition to the background assumption II being modally weak, the target assertion B is also modally weak. If rule (4) were applied, the syllogism would be paraphrased as:

*Since the Morning Star may be the Evening Star (II), and
since the Morning Star may rise (B), then
it is possible that the Evening Star may rise (B+II)*.*

Given the weak modalities of *both* the antecedents, the resulting clause should somehow be twice as "weak" as the clause resulting from substitution in A. Yet it is not clear what part of the assertion is "less likely" and how this should be represented (if at all). To eliminate such ill-formed inferences, we strengthen rule (4) to require that the modality of the BASS be *strictly*

⁶ Translations III(ii) and C(ii) are not very plausible for the given sentence, but they are included to demonstrate how the rules would work for such cases. They would be required in a sentence like "Necessarily the lion has a mane".

⁷ It may be that other facts which depend upon the truth of the background assumption can affect the acceptability of substitution, as in the following example:

*John may be the mayor, in which case Bill wouldn't like him.
Bill likes John
Therefore it is possible that Bill likes the mayor.**

weaker than the modality of the target assertion.⁸

5.4.2. Time of Descriptor Reference

Temporal constructs, and many modal constructs⁹, assert that descriptors within their scope refer at some time other than the time of utterance (*i.e.*, referred to as NOW). To substitute such descriptors, we must be sure that the equivalence relation between descriptors in the BASS includes the time of reference of the target descriptor.

For any sentence, two points of time must be determined: the time of descriptor reference, and the time of the (propositional) assertion. These are usually evaluated relative to the time of utterance of the target sentence. For example, in the sentence

(5-14) [A beautiful bride] will eat.

the assertional component "will eat" applies in the FUTURE; but it is unclear whether the descriptor "a beautiful bride" describes someone NOW or only in the FUTURE, at the time of the assertion.

The time of descriptor reference affects substitution according to the following general rule:

- (5) If the time of reference of the equivalent BASS descriptors overlaps the time of reference of the target descriptor, then the target descriptor can be replaced by the equivalent BASS descriptor. The descriptor in the resulting clause will have the original time of reference.

Consider example B:

⁸ Of course, this assumes there is a reasonable way to compare the relative strength of modalities. Certainly, "necessarily" is stronger than "possibly", but we also need to rank other modals such as "might", "may", "would", "certainly", etc. Even if there is a reasonable scheme, it will not be as straightforward as our rule (4) implies.

⁹ Many modals carry an implicit FUTURE operator (*e.g.*, *might* = ∇F). For example,
 (5-12) John *might* marry a beautiful woman.
 is equivalent to

(5-13) John *will possibly* marry a beautiful woman.

Thus, the time of the target descriptor reference becomes important as well as the modality of the proposition.

Background Assumptions:

I	<i>Mary is a beautiful bride.</i>	(def X: [beautiful_bride X]) <X = mary>
II	<i>Mary will be a beautiful bride.</i>	F (def X: [beautiful_bride X]) <X = mary>

Target Sentences:

A	<i>Mary is eating</i>	<eats mary>
B	<i>Mary will eat.</i>	F <eats mary>

Substitutional Inferences:

A + I	<i>A beautiful bride is eating.</i>	(def X: [beautiful_bride X]) <eat X>
A + II	<i>nil</i>	nil
B + I	<i>A beautiful bride will eat.</i>	(def X: [beautiful_bride X]) F <eat X>
B + II	<i>nil</i>	nil

Example B

Background assumption I asserts the current co-extensionality of the label "Mary" with the descriptor "beautiful bride", whereas assumption II asserts that sometime in the future the descriptors will be co-extensional. Targets A and B both refer to Mary NOW, but the assertion time of A is NOW while that of B is in the FUTURE.

The derivable inferences are fairly straightforward for the given example. The combination of target B with assumption II is somewhat interesting. One might be inclined to argue that, since labels always refer, the label "Mary" will have a referent at the time of the target assertion in B, namely in the future. Since assumption II provides a descriptor that is co-extensional in the future, the two should be intersubstituted. However, even though both descriptors apply at *some* future time in the future, this need not be the same point of time in the future. There is no way of knowing if the two future points overlap. Thus, our rule does not apply.

We have treated the aspect of time somewhat superficially. Our example includes only simple present and future. To implement the general rule with the multitude of tenses possible in NL (*e.g.*, simple past, progressive) requires a more elaborate model of temporal relations that is able to define more concisely when times "overlap".

5.4.3. Agents of Descriptors

In contexts with agents, even if the rules for intensionality and time are satisfied, substitution may fail if the descriptors to be substituted are not attributed to the same agent. In rules (1) to (5), the descriptor equivalences in the BASS were considered only from the listener's perspective. Now we must also consider what other agents (the listener believes) are also aware of the equivalence.

The general rule for substituting descriptors which are ambiguous with respect to descriptive content is this:

- (6) If the assertion of descriptor equivalence in the BASS is part of the knowledge base of the agent to 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.

Consider example C:

Background Assumptions:

- I *The fastest car in the world is Mike's Jaguar 300.* [fastest-car-in-world] = *r* = [mike's-jaguar-300])
- II *The fastest car in the world (always) is a Jaguar 300.* [fastest-car-in-world] = *a* = [jaguar-300]
- III *(Tom believes that) the fastest car in the world is Mike's Jaguar 300.* [fastest-car-in-world] = *t* = [mike's-jaguar-300]
- IV *(Tom believes that) the fastest car in the world is a Jaguar 300.* □ [fastest-car-in-world] = *t* = [jaguar-300]

Target Sentences:

- A *Tom bought the fastest car in the world.* (def X: [few X]) <bought tom, X>
- B *Tom wants the fastest car in the world.* (i) (def X: [few X]) <want tom, **X*>
 (ii) int_abs(C, def X: [few X]) <want tom, **X*>
 (iii) <want tom, *** (def X: [few X])>
 (iv) <want tom, **int_abs(C, def X: [few X])>*

Substitutional Inferences:

- A + I *Tom bought Mike's Jaguar 300.* (Rule 1) (def Y: [mj300 Y]) <bought tom, Y>
- A + II *Tom bought some Jaguar 300.* (Rule 1) (def Y: [j300 Y]) <bought tom, Y>
- B(i) + I *Tom wants Mike's Jaguar 300.* (Rule 1) (def Y: [mj300 Y]) <wants tom, **Y*>
- B(i) + II *Tom wants a Jaguar 300.* (Rule 1) (def Y: [j300 Y]) <wants tom, **Y*>
- B(ii) + I *nil* nil
- B(ii) + II *Tom wants a Jaguar 300.* (Rule 2) int_abs(C, def Y: [j300 Y]) <wants tom, **Y*>
- B(iii) + III *Tom wants Mike's Jaguar 300.* (Rule 1 & 6) <wants tom, **(def Y: [mj300 Y])>*
- B(iii) + IV *Tom wants a Jaguar 300.* (Rule 1 & 6) <wants tom, **(indef Y: [j300 Y])>*
- B(iv) + III *nil* nil
- B(iv) + IV *Tom wants some Jaguar 300.* (Rule 2 & 6) <wants tom, **int_abs(C, indef Y: [j300 Y])>*

Example C

Background assumption I asserts the co-extensionality of the descriptors "fastest car in the world" and "Mike's Jaguar 300", while assumption II asserts co-intensionality of the descriptors. Assumptions III and IV express the same equivalences,¹⁰ except knowledge of them is also attributed to the agent Tom.

Target sentence B is opaque with respect to descriptive content and intensionality of the descriptor "fastest car in the world". In translations B(i) and B(ii), the descriptor is attributed only to the speaker on extensional and intensional readings respectively. B(iii) and B(iv) are the same as (i) and (ii) except they attribute the descriptor to Tom. Target sentence A is a transparent version of target B.

To substitute in the target sentences A, B(i), and B(ii) we apply rules (1) and (2) for intensionality; no agents are involved. Translations B(iii) and B(iv) require applying rule (6) as well. B(iii) requires a co-extensional descriptor that Tom is aware of. Both BASS III and IV provide such a descriptor. B(iv) also requires a descriptor that Tom 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 in.

rules - p. 50

5.5. Extracting Descriptor Information from the MR

The MR, as is, contains all the information necessary to determine substitutable terms. However, it is not in a very convenient form for querying the knowledge-base. Usually the BASS will have already been assimilated into the general KB and will not be accessible in the form presented in the examples. In order to compare descriptors in the target to descriptors in the BASS, we extract the relevant aspects from our MR and express them explicitly by the use of *descriptor predicates*. A combination of these are then presented to the KB for comparison.

- (i) **desc**(*a*, *e*, *d1*) : The **desc** predicate ascribes a particular descriptor to an individual. The predicate is to be read "agent *a* would use the descriptor *d1* to describe the entity *e*".
- (ii) **label**(*a*, *e*, *name*) : The **label** predicate is used to indicate that the label *name* is known by agent *a* to be a label for the (individual) constant *e*.
- (iii) **time**(*t*, *e*, *d1*) : The **time** predicate asserts that descriptor *d1* is to describe entity *e* at time *t*.

The above predicates¹¹ are derived from the MR according to the following rules:

- (i) **Attributing descriptors to agents:**
Whenever a quantified term is used within the scope of the agent scope marker, its descriptor must be attributed to some set of agents using the **desc** predicate. In general, we have

$$\langle \text{Pred } a:\text{agent}, \dots, \wedge (\text{Det } X : [d1 \ X]) \dots \rangle \\ \rightarrow \text{desc}([a], X, d1).$$

For instance, from sentence B(iii) of example C, we derive **desc**(*tom*, *X*, fastest-car-in-world) and use this to query the knowledge base for co-extensional descriptors in Tom's knowledge base.

- (ii) **Attributing labels to agents:**
If a labelled constant is used within the scope of the agent opacity marker, a **label** predicate must be used to attribute that label to some set of agents. The derivation rule is similar to that of descriptors:

¹⁰ Instead of the full form, I have used the abbreviated descriptor relation notation for the BASS translations.

¹¹ There is no need to derive a predicate for intensionality since it is represented explicitly in the MR by the **int_abs** predicate.

$\langle \text{Pred } a:\text{agent}, \dots, \wedge \dots \text{label } c : \text{Name} \rangle \dots \rangle$
 $\rightarrow \text{label}([a], c, \text{Name}).$

Thus, from the sentence

(5-15) John wants Mary.

which has the representation " $\langle \text{want john}, \wedge (\text{label mary}, \text{Mary}) \rangle$ " we can derive $\text{label}(\text{john}, \text{mary}, \text{Mary})$ which asserts that John knows that "Mary" is *mary's* name.

(iii) Asserting the time of descriptor reference:

Whenever a descriptor is used within the scope of a future operator, a **time** predicate must be used to specify the future application of the descriptor. The rule for the future operator is:

$F \dots (\text{Det } X: [d1 \ X]) \dots$
 $\rightarrow \text{time}(\text{FUTURE}, X, d1).$

For example, from " $F (\text{def } X: [\text{beautiful-bride } X]) \langle \text{sings } X \rangle$ " we derive $\text{time}(\text{FUTURE}, X, \text{beautiful_bride})$. This can be extended to points of time other than just the FUTURE.

5.5.1. Knowing Labelled Referents and Substitution

If an agent knows the label of an entity, then we presume he knows of the referent as well, so that for an agent *a* and constant *c*

$\text{label}(a, c1, \text{Name}) \rightarrow \text{know_of}(a, c1)$

as in

$\text{label}(bf, bm, \text{"Brian Mulroney"}) \rightarrow \text{know_of}(bf, bm).$

But conversely, 'knowing of' the referent does not guarantee knowledge of a label for that referent. That is,

$\text{NOT } \Box (\text{know_of}(a, c1) \rightarrow \exists N \text{label}(a, c1, N))$

This is an important distinction for substitution. If the speaker knows the label of a specific object referred to and believes the listener does as well, then it is valid to substitute in the constant and label in place of the variable and descriptor. In terms of the above **desc** predicate this can be formulated as:

$\exists X \text{desc}(sp, X, d1) \text{ AND } \text{label}(sp, X, \text{name}) \text{ AND } \text{bel}(sp, \text{label}(l, X, \text{name}))$
 $\rightarrow \text{TARGET} = \text{TARGET } [d1/\text{name}]$

where *sp* is the speaker and *l* the listener.

5.5.2. Querying the Knowledge Base

Given the rules for establishing substitutable terms and the above descriptor predicates, we can now characterize the type of interactions with the KB that must occur to establish equivalent descriptors. (In the following capitalized letters represent variables, lowercase represent constants.)

(1) **Intensionality:**

The query posed to the KB to establish co-descriptors at the same level of intensionality should be:

Find a descriptor D1 such that either

- (a) $D1 = x = [\text{target_descriptor}]$ (co-extensions), or
- (b) $D1 = \wedge = [\text{target_descriptor}]$ (co-intensions), or
- (c) $D1 = d \neq \text{target_descriptor}$ (definitional equivalents).

The equivalence relation chosen should match the intensionality of the target

descriptor.

It is expected that the KB reasoner will be aware of these descriptor relations (or at least their corresponding long forms). Thus, the KB will be searched to find an appropriate descriptor D1. If the search is successful, the returned descriptor can be substituted in the target relation.

(2) Time of Reference:

To find a descriptor that refers at the "same time" as the target descriptor, we first extract the appropriate **time** predicate from the target MR. This is then presented to the KB to find a referent¹² for the descriptor at the relevant time. This is formulated as:

*Find all knowledge base entities, X, such that **time**(tt, X, [target_descriptor]) holds, where tt is the target descriptor time of reference.*

As a result, the KB should instantiate X to some KB entity, r. We must then query the KB a second time to search for other descriptors which describe the same entity:

*Find all descriptors D1 and their time of reference T for which **time**(T, r, D1) holds.*

If successful, this will instantiate D1 and T. If time T overlaps time tt, then it is valid to substitute D1 in the target clause.

(3) Agents of the Descriptors:

The procedure for determining co-descriptors with respect to agents is analogous to that for time. First, the **desc** predicate is derived from the target MR. This is used to determine a KB referent for the target descriptor:

*Find all KB entities, X, such that **desc**(a, X, [target_descriptor]) holds for agent a.*

Suppose X gets instantiated to the KB referent r. We then must find any other descriptors the agent would use for the same type of referent:

*Find all descriptors D1 such that **desc**(a, r, D1) holds.*

Any instantiations found for D1 can then be substituted in the target sentence. To find labels that could be substituted, we would use a comparable request:

*Find all labels L such that **label**(a, r, L) holds.*

Then the instantiated value of L could be substituted for the target descriptor.

¹² This could be an individual, an individual concept, or a universal concept.

CHAPTER 6

NLP Systems and Opacity

Four factors must be considered if a NLP system is to handle linguistic opacity adequately:

- (1) recognition of opaque contexts,
(SYNTAX plus semantics)
- (2) representation of ambiguities,
(MEANING REPRESENTATION)
- (3) interpretation and disambiguation of intensional operators,
(SEMANTIC INTERPRETATION plus pragmatics)
- (4) determination of valid inferences.
(PRAGMATICS)

Linguists have typically addressed issues (1) and (2), while logicians tend to focus more on (3) and (4). Our main emphasis is on (2) and (3). Of course components (1) and (4) can not be overlooked; they are integral parts of a complete system. In chapter 2, we outlined the syntactic features of opaque constructs. This enables us to construct a set of rules to recognize and parse opaque contexts; no claim however is made that they are the optimal parsing rules. On the other hand, the proposed semantic classification is based on which inference rules were thought to fail in the various contexts. Thus, the relevant aspects of (1) and (4) have been given consideration. The remainder of this chapter describes the implementation of a meaning representation for opaque constructs and suggestions for its interpretation in a knowledge-based system.

6.1. Computability of Logic Grammars and Montague Semantics

Several of the fundamental assumptions characterizing Montague's (1973) semantic system can be applied to the construction of a NLP system. Montague's system is appealing since:

- it is based in logic (an intensional logic),
- it is fully compositional both in syntax and semantics,
- it has a one-to-one correspondence between syntax and semantic rules and objects,
- it provides a uniform, systematic treatment of opacity,
- it has a well-developed method for interpretation with respect to a chosen model.

Montague's "PTQ" system is not directly applicable in AI contexts. It must be modified so that it is not essentially *truth-conditional* (Hirst 1986). While one property of a sentence may be that it can or does refer to a truth value, it will have many other attributes which together determine its appropriate "meaning". Also, to be computationally tractable, the system must use finite objects rather than unwieldy infinite sets and total functions. And, instead of a model-theoretic interpretation, we will need to formulate a comparable knowledge-based interpretation suitable for NLP.

Our modified AI-version of Montague semantics is realized with a logic grammar. Logic grammars are a parsing formalism available in most logic programming environments (*e.g.*, Prolog), and are well-suited for NLP (see Warren and Pereira 1980). Since the unification

operation in Prolog matches arbitrary data structures, it seems that logic grammars can support higher-order quantification and lambda abstraction¹. As well, since partial information and partial entities are easily manipulated in Prolog (using logic variables and unification), abstract items such as intensions can be represented.

Montague-like semantics have been implemented previously by several different groups. Other parsing mechanisms comparable to those based on logic grammars have been coupled with Montague semantics, notably the Phrase Structure Grammar of Gazdar (Schubert and Pelletier 1982) and an ATN-driven approach (Friedman and Warren 1978). Hirst (1986) has developed these ideas for conventional AI-NLP systems. However, none to date have dealt explicitly with opaque constructs and none within the logic programming framework.

6.2. Specifications for an Implementation

To demonstrate the potential of our MR in a computational environment, a moderate-size NLP interface has been developed. The component has been designed in three stages to facilitate later expansion and support from a knowledge-based AI system. The goals of each of these stages are as follows.

(1) Parsing into an initially ambiguous logical form:

Following Montague, we have two sets of translation rules: one for syntax and one for semantics, applied in one-to-one correspondence. The syntax rules check the well-formedness constraints of the NL and construct a constituent tree; the semantic rules translate the constituent tree structure into a formula of our MR.

Montague adheres strictly to the one-to-one correspondence of rules, and requires that his IL formulas be both structurally and semantically unambiguous. To ensure this, his system includes syntax rules that are unmotivated by the surface structure but necessitated only by the semantics (*e.g.*, rules S14 and S16 which determine the various scopings of quantifiers). Instead, I have chosen to parse into a *structurally* unambiguous but semantically ambiguous MR. For a sentence fragment like "wants a coat like Bill's" that can have four interpretations, the parser initially produces a single MR formula, but indicates which readings can be derived from it. Computationally it make sense only to parse what is there and not invent new syntax rules just to get the semantic rules matching. In NLP systems, this one-to-one correspondence is motivated by practical issues (*e.g.*, early semantic filtering, modularity) and not ideological principles, so it is reasonable to relax the constraint somewhat to permit more efficient parsing. Furthermore, at this point the parser doesn't have sufficient information to make a prudent selection as to the most appropriate reading. Thus we have chosen to generate several separate MR formulas when one will do until more complete information is available.

(2) Expanding the initial formula:

The second stage expands the formula from the initial parse to yield a set of plausible readings for the input. During this phase, the parser may appeal to the *immediate pragmatic context* to resolve local referents, indexicals, and to constrain the possible type of referents for descriptions. The immediate pragmatic context includes representations of the *discourse situation* and the *discourse focus*. The discourse situation is a tuple containing KB identifiers for the current speaker, listener, time, and location. These are intended as points of reference to help resolve (explicit) indexicals of time and place, and provide "handles" into the KB for more detailed information about each. The discourse focus is a list of entities (either intensional or extensional) which have been "recently" referred to in the discourse. Each entry is a (*KB id*, *descriptor*) pair. The KB id is a key to access the referent's KB representation, and the

¹ D.H.D. Warren (1982) has considered how higher-order variables and lambda expressions might be incorporated into standard Prolog framework.

descriptor gives the most recent means of describing it. At the completion of this stage, we have a MR formula as defined in chapter 4, ready to be presented to a KB component for further processing (*e.g.*, discourse understanding, database querying).²

(3) Generating valid inferences:

Once the parse is complete, the MR³ is passed to the *pragmatic post-processor* which provides the main interface to the general KB.⁴ It is responsible for extracting and applying KB information to solve any unresolved referents and to derive linguistic inferences. Thus, it generates the *desc*, *label*, and *time* predicates necessary to query the KB about substitution and other inference operations. As well, the pragmatic post-processor initiates KR processes to assimilate new information and new descriptors from the MR into the KB. For these purposes, it maintains the *discourse environment*, an elaboration of the discourse situation to include information about the extralinguistic setting and a representation of the previous discourse.⁵ The discourse environment need not be a separate data structure but can consist of a set of procedures for accessing the relevant information in the general KB. The pragmatic post-processor also keeps track of the current and compatible states of affairs. In particular, it must update the priority and relevance of the compatibility constraints which define the states of affairs.

Figure 4 depicts a standard NLP system and indicates which aspects are part of our implementation. Those marked with an asterisk are **not** included, while those marked with a plus sign are given emphasis. In the parser we have not taken account of structural or lexical ambiguities unless it was necessary to distinguish among different opaque readings. We have implemented only those parts of the pragmatic post-processor which demonstrate how the rules for substitution are restricted in opaquely marked contexts. That is, it deduces from the MR the descriptor predicates necessary for establishing substitutable terms. To show how the descriptor predicates relate to a KB component, we include some arbitrary information to simulate what might be available from a KB, but otherwise do not include the KB-related aspects of the post-processor.

6.3. Refinements of the MR for the NLP System

The current implementation is in C-Prolog 1.4 (Pereira 1984) on a VAX-11/780. The grammar rules are specified with SAUMER, a logic grammar parser developed at Simon Fraser University by Fred Popowich (1984, 1985). SAUMER basically supports a generalized phrase structure grammar. A typical rule is of the form " $\alpha \rightarrow \beta : \gamma$ " where α and β represent the left-hand and right-hand sides of the syntactic re-write rule, and γ encodes the corresponding semantic formula. α and β are a sequence of terminals and non-terminals, where each non-terminal can take a finite number of terms as arguments. β may also contain calls to arbitrary Prolog predicates.

6.4. Types of Objects in the Implemented MR

To be compatible with the SAUMER parser and Prolog in general, the implemented MR has been syntactically modified as outlined below.

² If the various readings are required earlier, it is straightforward to expand the MR formula before the parse is complete for that sentence.

³ At this point one or two of the candidates should have been selected as the most likely; this saves processing them all.

⁴ This is not the only possible scenario, but it is sufficient to illustrate how the MR formula can be used.

⁵ This is somewhat like a short-term memory exclusively for information pertaining to the discourse.

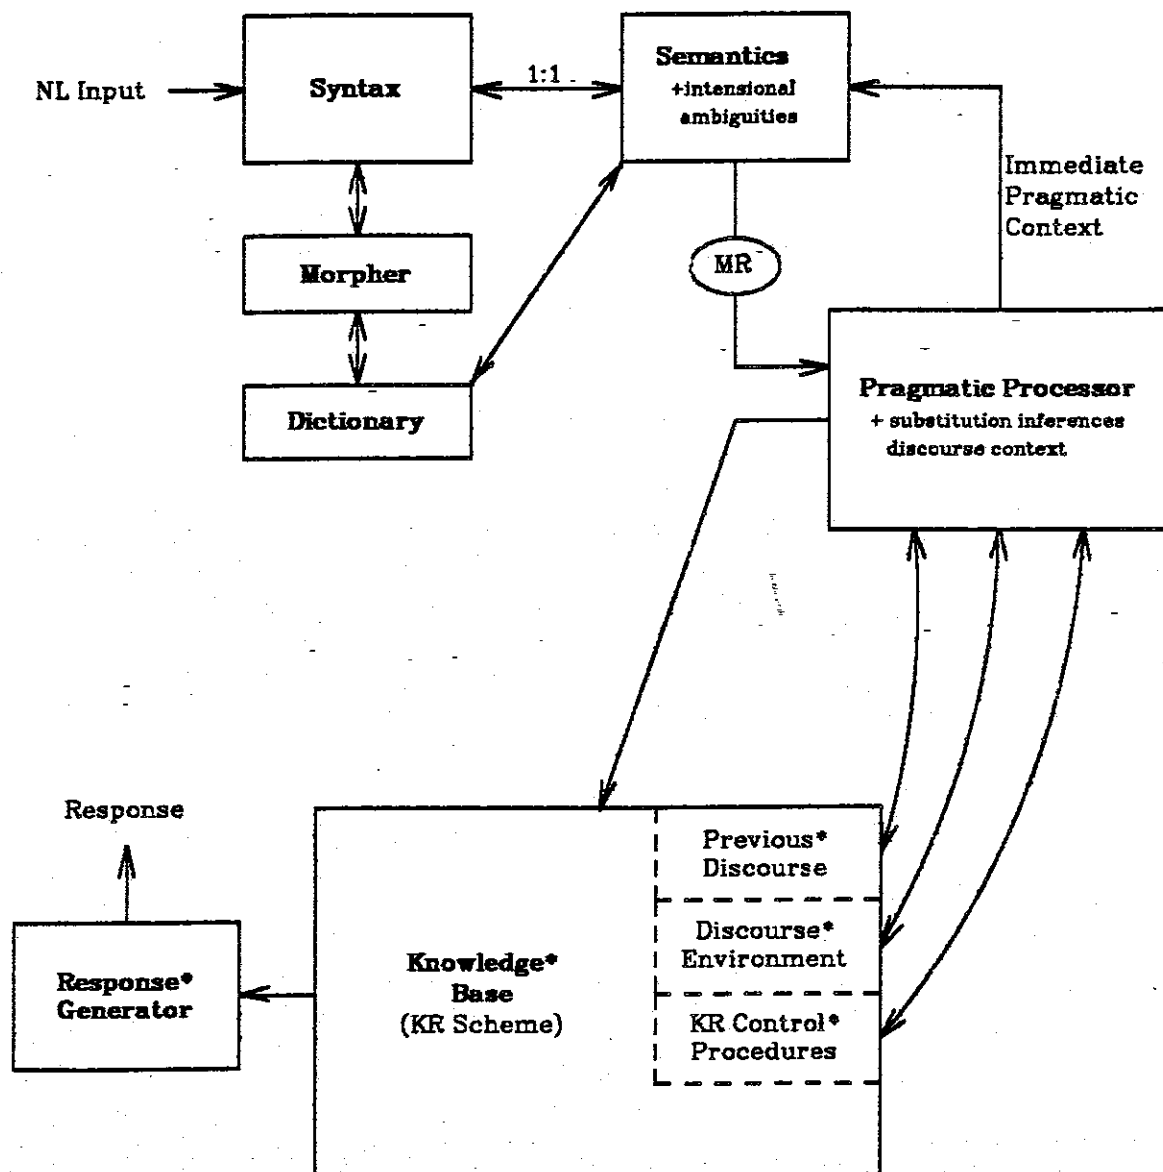


Figure 4. Typical NLP System

(1) Terms:

Constant terms such as *john1* are a sequence of lowercase letters followed by a numerical suffix. Unknown referents are represented by variable terms of the form *varN* where *var* belongs to {*x*, *y*, *z*} and *N* is a numerical suffix. (Note that these are not Prolog variables.) Quantified terms have the general form (*determiner, variable, restriction*) as in (*some, x1, man2*).

(2) Predicates:

Predicate application is the primary means of combining constituents. It takes two forms:

- (i) $[P, \text{Arg0}, \text{Arg1}, \dots, \text{ArgN}]$: The predicate *P* is applied to the list of arguments (*Arg0*, *Arg1*, ..., *ArgN*) and the result is another predicate,
- (ii) $\langle \text{Arg0}, P, \text{Arg1}, \dots, \text{ArgN} \rangle$: The predicate *P* is applied to the list of arguments (*Arg0*, *Arg1*, ..., *ArgN*) and the result is a proposition. It is similar to $[P, \text{Arg0}, \dots, \text{ArgN}]$ except the predicate is in second position and the result cannot be further applied, except to be evaluated in the knowledge base.

Predicate application is implemented using the idea of leftmost section (Schubert and Pelletier 1982) so that $P(Y)$ is equivalent to $[P\ Y]$. Thus, $[[\text{give}, \text{mary}], (\text{a } y3 \text{ dog4})]$ yields $[\text{give}, \text{mary}, (\text{a } y3 \text{ dog})]$, and, at the sentential level, $[[\text{give}, \text{mary2}, (\text{a3 } y4 \text{ dog5})], \text{john1}]$ results in $\langle \text{john1}, \text{give}, \text{mary2}, (\text{a3 } y4 \text{ dog5}) \rangle$.

(3) Other Operators and Predicates:

Lambda abstraction takes the form of $x \text{ lmda } Formula$ where *x* is any variable occurring one or more times in *Formula*. When applied to a single argument, that argument is substituted for all occurrences of *x* in *Formula*. However, *Formula* may not contain nested lambda expressions using the same variable.

In the implemented MR, logical *and* and *or* become **and** and **or** respectively, while material implication becomes **cond**.

6.5. The Knowledge Base Interpretation

A meaning representation that distinguishes levels of intensionality and which attributes descriptors to agents does more than contribute to linguistic completeness; such information can guide the formulation of KB expectations and dictate successful reasoning strategies. For example, a specific reading of a NP description requires that the KB contain (or should create) a corresponding instance, whereas on a non-specific reading the NP specifies properties to identify a class.

Once the type of referent has been established, expectations can be set up in the KB for subsequent references using a similar descriptor. Nick Cercone and Tomek Strzalkowski (1985) have formulated rules for short (*i.e.*, two-line) stories which disambiguate intersentential references made in opaque and regular contexts.⁶

Likewise, the ascription of descriptors to agents involved in the discourse is crucial to the building of accurate user models. Such models, often called *belief spaces*, have assumed an increasingly significant role for both discourse understanding (*e.g.*, speech acts) and more general reasoning processes in knowledge representation.

6.5.1. Adequacy of Available KR Schemes

Opaque contexts and their representations poke at the boundaries of the expressive ability of traditional frame-based KR schemes. We characterize what facilities are necessary for an

⁶ In their terminology, such constructs are termed *imperfect* and *perfect*.

adequate KB interpretation of our MR. To do this, several knowledge representation schemes are referred to, including KL-One (Schmolze and Brachman 1982), PSN (Levesque and Mylopoulos 1979), KRL (Bobrow and Winograd 1977), Prolog/KR (Nakashima 1983), and EFR (Nishida 1983).⁷

The typical frame-based KR scheme allows the definition of two types of entities: *generic* frames (*i.e.*, CLASSES) and *individual* frames (*i.e.*, INSTANCES of classes). A frame of either type is a collection of properties and their values (commonly termed SLOTS and FILLERS) which together constitute a logical unit. Many KR languages (*e.g.*, KL-One⁸, KRL) construe such a frame unit as a description which can be applied to objects. The basic strategy in the context of NLP is to match linguistic descriptors from the input against those in the KB to determine an appropriate referent, usually an instance of the described class. The KB is then updated to reflect new information being asserted. Thus, for the sentence

(6-1) Sue bought [a coat].

the KB matcher would find the instance node for "Sue" (of class PERSON), create an instance node of class COAT, and finally assert an instance of the relation BUYING between the two nodes. Applying the same strategy to the opaque sentence

(6-2) Sue *wants* [a coat].

would, however, erroneously assert that Sue wanted a *particular* instance of a coat. Thus, things are rather different for interpreting opaque contexts.

Minimally we require a KR scheme which distinguishes different levels of intensionality and can assert properties and relations at each level (especially equivalence relations). Additionally, the knowledge must be partitioned according to agents, time, and hypothetical states of affairs so that information for substitution and inference can be efficiently extracted. We discuss each in turn.

6.5.2. Intensionality

The first requirement of the KR scheme is that it be able to represent individual and universal concepts as well as (individual) instances of them:

In designing a network to handle intensional entities, we need to provide for definite entities that are intended to correspond to particular entities in the real world, indefinite entities which do not necessarily have corresponding entities in the real world, and definite and indefinite variable entities which stand in some relation to other entities and whose instantiations will depend on the instantiation of those other entities. (Woods 1975 : 68)

Woods's "definite entities" are what I term (extensional) individuals, while his "indefinite entities" correspond to either universal or individual concepts. Each of these notions must be given appropriate analogs in the KB. Mechanisms also must be available to introduce novel concepts, to combine concepts to form larger, non-atomic concepts, and to modify concepts on the basis of new information and/or experience.

Standard KR schemes lack the ability to refer to conceptual nodes apart from defining and using them as descriptions. In example (6-2) above, the descriptor "a coat" applies not to a particular individual, but to the (abstract) concept of an individual coat. Some languages such as

⁷ None of these are described in detail; the reader is referred to the relevant references for further explanation.

⁸ Strictly speaking KL-One is a semantic network approach; however frame and semantic networks are syntactic variants of the same style of representation. Both associate nodes with sets of properties and impose an organizational structure on them which establishes their core meaning and regulates the inheritance of properties.

KL-One refer to generic frames as *concepts*. These frames embody the properties of the corresponding concept, but the frames themselves cannot be referred to as concepts, nor do they apply to objects which are concepts. Rather they group together abstract properties which describe extensional instances. Because these generic descriptions cannot be referents of other frame descriptions, we are unable to make assertions about them; they cannot fill the slots of individual instances; only other instances can. This issue is briefly addressed in the *Summary of the KL-One Language* under the guise of "meta-descriptions":

In addition to the use of Nexuses as "surrogates" for outside entities, KL-One allows reference to internal entities (*e.g.*, Concepts) as well. Thus, one can "meta-describe" a KL-One object in KL-One...These are not yet part of the implemented system. (Schmolze and Brachman 1982 : 257)

Work on KR schemes with specifically intensional nodes has been modest (see Maida and Shapiro 1982, Creary 1979). Some critical issues to decide are:

- (i) whether two distinct nodes are needed to represent every object, one for its intension and one for its extension;
- (ii) what kind of representation is appropriate for an intension (*e.g.*, concept node, function, description, set of procedures) and whether it should be explicit or implicit until required;
- (iii) whether every entity can have an intension. If this is enforced, it is possible to have an infinite hierarchy of intensions of intensions of intensions and so on.
- (iv) how to assert knowing *whether* a description or sentence applies, as distinct from knowing *that* a description or sentence applies.

Relative to our work, these decisions affect how the intensional abstraction function, *int_abs*, is defined. Usually, when searching for KB referents, a quantified term provides a description to match against those already in the KB under the assumption that they apply to some extensional entity. However, if a quantified term is the argument of an *int_abs*, then its description is used to search for a conceptual node as a referent instead of an instance. (If new information is being added, the description would be used to provide definitional properties for a new conceptual node.)

Universal concepts can be represented by generic frames considered independently of any instances. In KL-One this corresponds to a GENERIC CONCEPT; in KRL this is a PROTOTYPE unit; and in PSN a CLASS type would be appropriate. The aggregate of its properties (ROLES in KL-One, DESCRIPTORS in KRL, SLOTS in PSN) define the attributes that any member should have, and may also include properties of the class itself (*e.g.*, cardinality).⁹ For example, the sentence

(6-3) [A unicorn] has green eyes.

asserts that all members of the class UNICORN must have a slot EYE-COLOUR with the filler GREEN.

Individuals can be interpreted as *individual* frames to stand for objects which exist in *some* state of affairs. These instances are fully specified frames which have values for all their defining slots. It is reasonable to ask in which states of affairs such individuals "exist" and what descriptions apply to them in each of these states. To interpret proper nouns, we assume that all

⁹ Many KR languages use slots for default properties whose values can be overridden to handle exceptions. In this case the frame does not actually define anything but rather represents a stereotypical member (see Brachman 1983). To preserve the definitional interpretation of the frames, we assume non-cancellable properties. It would seem that to handle defaults, an intermediate class of *prototypical* properties should be permitted as part of the conceptual taxonomy.

individuals will have a slot to contain an individual's linguistic label (if known). This idea of an individual is comparable to INDIVIDUAL units of KRL and INSTANCE nodes of PSN. In KL-One NEXUSES denote individuals, but not as frame structures. Rather they are simple reference points with no internal structure; information about them is exclusively through the descriptions applied to them or the assertions they participate in.

Individual concepts are more difficult to find a suitable KB interpretation for. An individual concept does not correspond to a particular individual but it is presumed that such a unique referent is possible in a state of affairs compatible with the current one.¹⁰ In example (6-2), the assertion is not about the concept of *any* coat but about the concept of some unknown but *individual* coat. The individual concept specifies those properties which are currently known about an individual. Since it is possible that more may be learned, the KR scheme must allow further specialization of individual concepts. An individual concept may uniquely identify individuals only in certain states of affairs. Therefore, its definition should include restrictions for the set of contexts in which it is applicable. In a traditional CLASS frame, an individual concept intuitively corresponds to an INSTANCE-OF slot, as Figure 5 indicates. This reflects the idea that it describes an individual instance of a class but with an unknown referent.

KRL provides a close approximation to an individual concept with its notion of a "ghost":

A representation must enable us to describe entities whose unique identity is unknown. There are many cases in which we may know many properties of some [individual] object without knowing which of the known objects in our world it is. (Bobrow and Winograd 1977 : 12)

For this, KRL defines a MANIFESTATION unit, which groups together a set of descriptions intended to apply to *some* individual. This accentuates the difference between knowing *of* an individual (*cf.* individual concept) without knowing *who* that individual is (*cf.* individual).


Nishida's (1983) solution for *intensional expressions*¹¹ in EFR is to permit indefinite, definite, or set quantifier nodes to introduce variables over instance nodes. If a referent is subsequently found satisfying the asserted intensional properties, the variable is "extensionalized" by having a constant node take its place. After this extensionalization there is no longer a conceptual representation of the individual. Thus no further reference to or reasoning about that concept is possible. The only way to indicate relations among such concepts is through the instantiated referents. This makes it difficult to assert co-intensionality between two individual concepts.

KL-One includes a node type termed an *individual concept*. It is a description which can apply to at most one individual in each context. Apparently the only feature distinguishing it from a generic concept is that its referent must be unique. If it is truly an individual concept description (and not the description of one particular known object)¹², then it can model our notion of an individual concept. As part of its definition we would add the expectation that it be applicable to some (unknown) object in a compatible state of affairs. For this we would introduce variables over individual nodes and quantify them relative to the states of affairs they may occur in.

¹⁰ The current state of affairs is the state of affairs in which the description is introduced. It is up to the KB reasoner to determine such things as uniqueness of reference, co-extensionality of descriptors, etc. in a particular state of affairs.

¹¹ "Those [expressions] referring to objects whose existence is not assured" (Nishida 1983 : 32).

¹² In some examples it seems that KL-One's individual concept nodes are used as elaborate INSTANCE nodes, for example in Schmolze and Brachman (1982 : 240), a particular arch is described by an individual concept node labelled ARCH#1 with supports SQUARE_BLOCK#1 and SQUARE_BLOCK#2. If an individual concept is like an instance, then is the same as our INDIVIDUAL and is of little help, since it is intimately connected with a known entity.



- (1) [A unicorn] is a quaint animal. (Universal Concept)
- (2) Winston seeks [a unicorn with green hair]. (Individual Concept)
- (3) Sylvester met [a unicorn with pink eyes]. (Individual)

Figure 5. Intensionality of Referents

6.5.2.1. Asserting Relations Among Such Nodes

Representing and recognizing the various types of descriptor relations in the knowledge base is essential for carrying out substitution and other reasoning procedures. For each type of copular link discussed in Chapter 4, I suggest an approximate KB representation. It is difficult to choose the correct representation; often several are possible for a given KR scheme with some being easier to modify, some easier to do matching and retrieval with, and some easier to reason about.

For the examples, I adopt a KL-One notation and intend that the semantics be that accepted by the general KL-One community. To avoid confusion I will define the KL-One entities of interest and how I interpret each.

"A KL-One network consists of two distinct parts: a taxonomy of intensional descriptions [Concepts] and an extensional database of objects [Nexuses] to which descriptions can be attached" (Lipkis 1982 : 130). *Concept* nodes (drawn as ellipses) are a collection of definitional properties, termed *roles*, which denote the corresponding concept. These concepts are considered as descriptions which can be applied to extensional objects. However, describing a concept is not always the

same as denoting a concept. If there are several ways to *describe* the same concept¹³, then there can be multiple node structures in the KB; but if a node structure *denotes* the concept, then there should only ever be one such structure in the KB. In theory, KL-One subscribes to the latter idea — there is only one node per concept and its properties are used to access instances of that concept. In practice, however, it is computationally difficult to decide that two descriptions embody the same concept. Some type of (generalized) unification algorithm must decide if and how two conceptual representations coincide. KL-One relies exclusively on structural comparison and subsumption relations for this.

A generalization/specialization hierarchy among concepts is maintained using a *subsumption cable* (solid line arrow). Whenever a concept description applies, so do all of its subsuming concepts. *Generic* concepts may describe several objects in one context while *individual* concepts describe at most one individual. Individual concepts are linked to their subsuming generic concept by an *individuation cable* (dashed arrow), which I interpret this as an individual abstraction link and not as an instance link.

Concepts are used to describe and make assertions about entities, called *nexuses*, within a particular state of affairs, a *context*. This is asserted by using a *description wire* (squiggly line) to connect the describing node to a nexus. Strictly speaking, nexuses correspond to individuals only, but we would like them to range over individual and universal concepts as well. KL-One permits contexts to be referred to (using “meta-anchors”) but little work has been done to provide primitives for structuring and manipulating them.

When interacting with a KL-One knowledge base, it is important to maintain the distinction between assertional and definitional knowledge. *Definitional* information provides a set of context-independent conceptual descriptions that can be used to identify entities. *Assertional* information deals with incidental properties and relationships between objects in particular states of affairs. Considered in our framework, the conceptual taxonomy provides definitional knowledge which is true over all compatible states of affairs, while asserted knowledge is more volatile and holds only in a limited set of states of affairs. As a result, to make assertions about concepts in non-compatible states of affairs would require another version (or a variation of) the definitional taxonomy.

¹³ For instance the concept of a unmarried male person could be retrieved using any of the descriptors “bachelor”, “unmarried man”, “single male”, or “human of the male gender who does not have a wife”.

6.5.2.1.1. Ordinary Classification

In ordinary classification, an individual instance is asserted to be a member of another class. That is, some generic concept (already existing in the KB) is applied to further describe a known individual in the current state of affairs.

The Prime Minister is a Conservative.

(def X : [prime_minister X]) <conservative X>

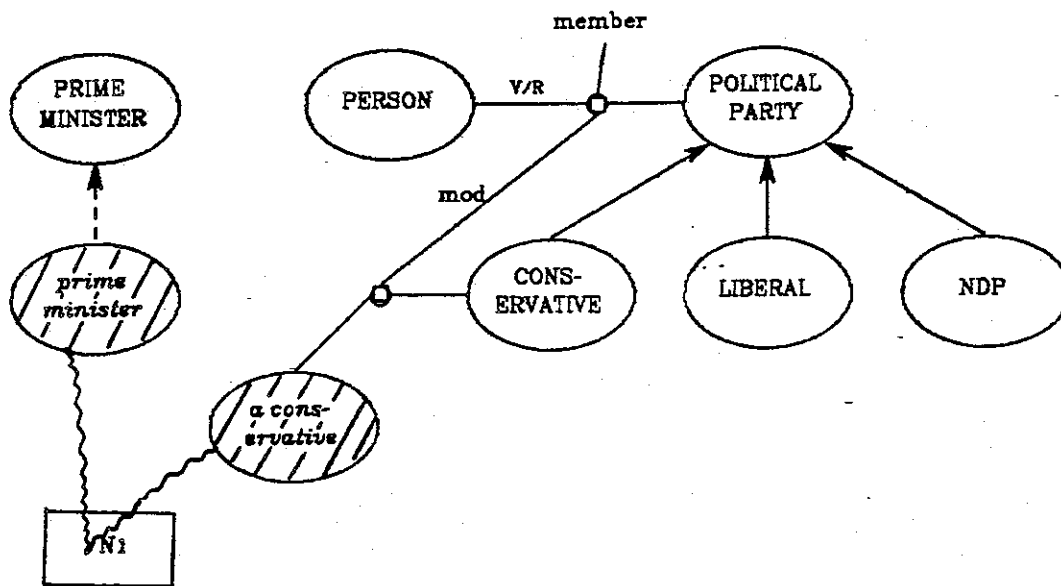


Figure 6. Ordinary Classification

6.5.2.1.2. Generic Classification

Generic classification indicates that one concept is a more specialized sub-concept of another concept. The child concept inherits all the properties of the parent plus adding or refining some to particularize it. The equivalent of an ISA link under the interpretation of "conceptual containment" (Brachman 1983) is needed to represent this. KL-One's version is the subsumption cable which indicates that whenever the description of the child concept applies so must the parent description.

The Prime Minister of Canada is an elected official.

`int_abs(C, def X : [pm_of_canada X]) <elected_official C>`




Figure 7. Generic Classification

6.5.2.1.3. Extensional Equivalence

Extensional equivalence asserts that two concept nodes describe the same referent in a particular context. Thus in KL-One, a descriptor wire should be connected from each of the two concept nodes to a single nexus. The KB reasoner must check every time a new description is applied to see if it co-refers with any other descriptions. Often the KB may discover that two entities previously thought to be distinct are in fact the same in which case the two individual nodes should be merged. *Realization* algorithms to carry out such merge operations have been designed and implemented for KL-One (Mark 1982).¹⁴ They must however be carefully constrained to ensure computational tractability.

¹⁴ The converse operation of node splitting may also be required if it is subsequently found that two concepts thought to describe the same object in fact describe two distinct objects.

The Prime Minister is the President of Alcoholics Anonymous.

(def X : [prime_minister X]) (def Y : [president_AA Y]) <X = Y>

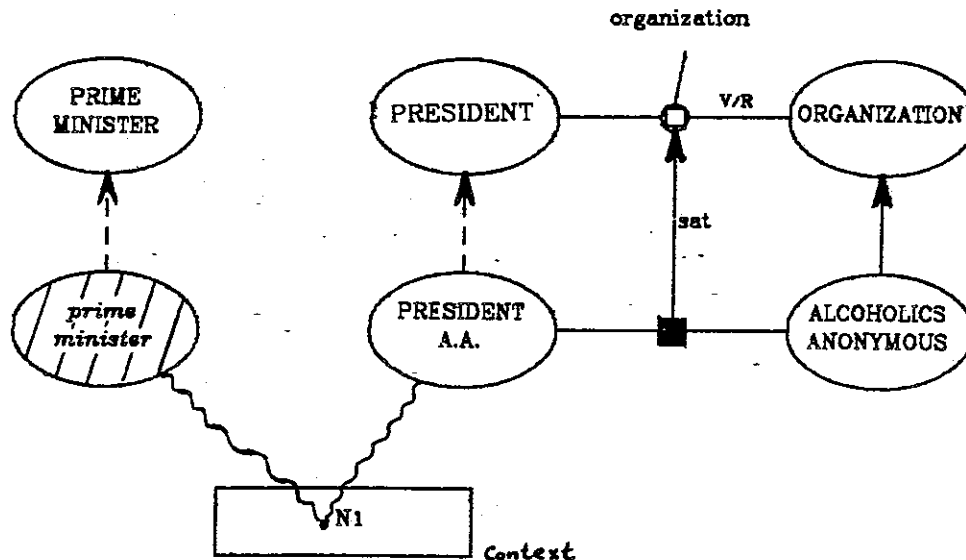


Figure 8. Extensional Equivalence

6.5.2.1.4. Intensional Equivalence

Intensional equivalence between two nodes specifies that they must co-refer across all compatible states of affairs. That is, even though they express distinct concepts, the two nodes will always apply to the same referents. If two descriptions are completely co-intensional, then they must co-refer over all times and over all states of affairs. To do this in KL-One requires using *role value maps* to stipulate that the roles which the co-intensional nodes fill must always describe the same extension but without asserting definitional equivalence.¹⁵

¹⁵ Another possibility is to use a bidirectional subsumption cable between the two concept nodes. However, this link cannot be the same as that used in generic classification. Instead of asserting conceptual containment, it should assert the weaker notion of *description* containment. The assertional capabilities of the concepts are being equated, not their meanings.

The Queen of England is always the commander of the Navy.

$\text{int_abs}(C1, \text{def } X : [\text{queen_Eng } X]) \text{ int_abs}(C2, \text{def } Y : [\text{cmd_navy } Y]) \square \langle X = Y \rangle$

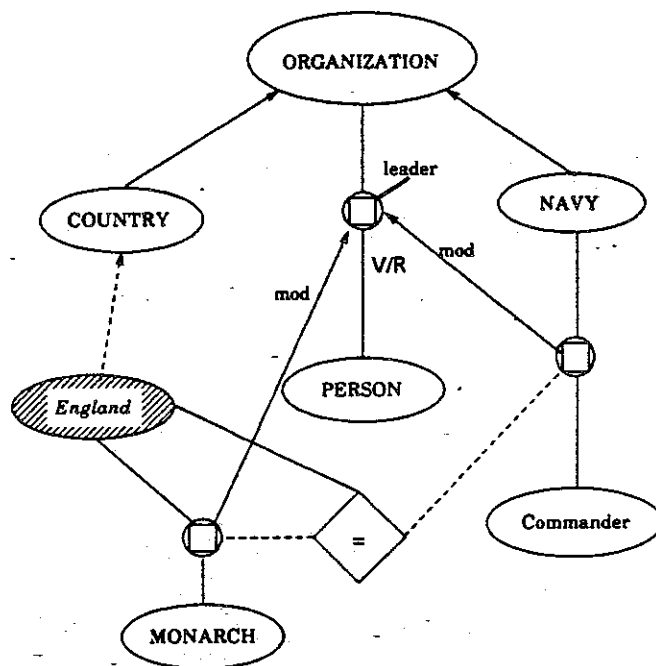


Figure 9. Intensional Equivalence

6.5.2.1.5. Definitional Equivalence

Definitional equivalence may equate two atomic descriptors as being the same concept (*i.e.*, single word synonyms), or it may assert the equivalence of composite descriptors (*e.g.*, "A picador is the horseman who incites the bull in a bullfight by pricking its neck with a lance."). In the first case, the synonymous descriptors must simply be mapped onto the same conceptual structure in the KB; there is only one structure, but two linguistic phrases naming it as Figure 10a illustrates.

groundhog =df woodchuck




Figure 10a. Definitional Equivalence of Atomic Descriptors

Composite descriptors may lead to several structures representing the same concept, but each with a different internal structure. In a perfect world, one should never need to actually represent that two concepts are equal. There should be only one place in the taxonomy per concept. When two or more descriptions which describe the same "virtual" concept are created as distinct structures, this should be recognized and the two merged. However, this is not always possible (and perhaps not desirable) due to computational limitations. It is not clear how to optimally unify such distinct but equivalent descriptions. (e.g., The KL-One reasoner checks only for structurally identical concepts.) Instead of merging, such equal concepts can be marked as identical siblings in the definitional taxonomy using a mutual subsumption cable between them.¹⁶ Figure 10b depicts two different structures for the concept "bachelor". Structurally they are distinct but semantically they should be equivalent.

Aside from recognizing the two as equivalent, there is another problem with the individual structures in Figure 10b. If a "bachelor" is truly definitionally equivalent with "unmarried man", the two should be interchangeable in all contexts.¹⁷ However, the subsumption arrow points only in one direction, which can be read *Whenever "bachelor" describes an entity, then so does "unmarried", "male", and "person"*. However, it does not say that *whenever "unmarried", "male", and "person" apply, so must "bachelor"*; this deduction would have to be drawn by a process external to the KB.

¹⁶ This is not well-defined in KL-One since it leads to circular subsumption, which is problematic for the KR classifier.

¹⁷ It is questionable in this case if the two are equivalent. "Bachelor" seems to carry a few connotations which are not part of "unmarried man", such as "eligible" and "heterosexual" (Winograd 1976). Perhaps we should use "swinging unmarried man" as its definitional equivalent.

A bachelor is an unmarried man.

$\text{int_abs}(C1, \text{indef } X : [\text{bachelor } X]) \text{ int_abs}(C2, \text{indef } Y : [\text{unmarried_man } Y]) < C1 = C2 >$

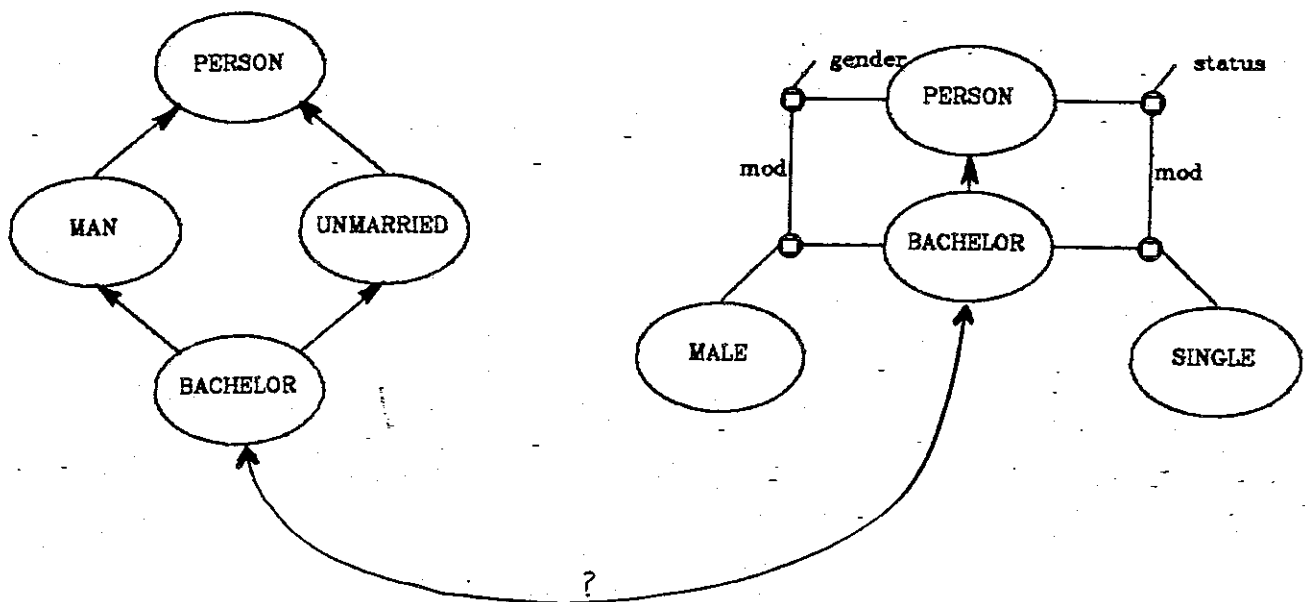


Figure 10b. Definitional Equivalence

6.5.2.1.6. Extensional Predication

Extensional predication resembles ordinary classification, except that the classifying node designates a property rather than a set of objects. In most cases these are two ways of looking at the same phenomenon. By creating a class of objects with a common property, checking for class membership is equivalent to checking that an object has that property. There is a choice whether to represent it as class membership or as predication. Both should be available in the KR scheme, and a decision made at interpretation time as to which is more appropriate.¹⁸

¹⁸ In KR languages in which frames are used for individuals, predication would be represented by filling an appropriate slot in the frame with the predicating information. Class membership would be asserted by an instance link from the individual frame to the classifying generic frame.

The (current) prime minister is a fool.

(def X : [prime_minister X]) <fool X>




Figure 11. Extensional Predication

6.5.2.1.7. Intensional Predication

Intensional predication requires that a new property be associated with a concept as part of its definitional (or prototypical) knowledge. Thus a new property is attached to the conceptual node.

The Prime Minister is (always) well paid.

$\text{int_abs}(C, \text{def } X : [\text{prime_minister } X]) <\text{well_paid } C>$

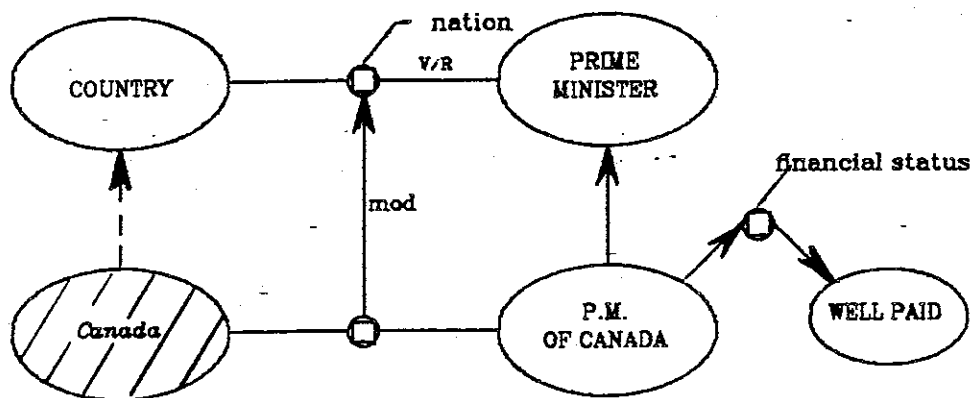
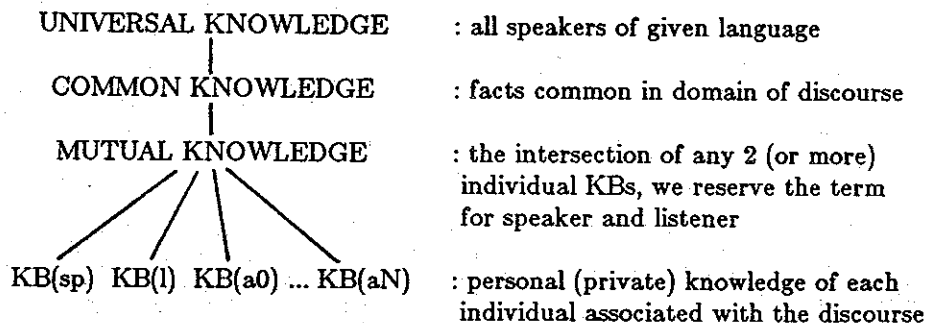


Figure 12. Intensional Predication

6.5.3. Agents in the Knowledge Base

The knowledge base is assumed to be composed of many interdependent *virtual* KBs organized according to who is presumed to have access to the knowledge contained therein (*cf. belief and want spaces of Allen (1979)*).



Since we can only know things from one individual's outlook, we view all knowledge from the listener's perspective. All the axioms of our complete KB would be implicitly prefixed by "the listener believes that", perhaps a coloured view of the world but nonetheless useful.

To evaluate descriptors, the KB reasoner may have to take into account the *belief* space to which the description is ascribed. This information is made explicit from our MR with the *desc*

predicate. Depending on which of its arguments are instantiated, $desc(a,o,d)$ can be used to extract three types of information from the KB:

- (1) what agents would describe object o by description d ,
- (2) what object agent a would describe by d ,
- (3) what descriptions agent a would use to describe object o .

Such a partition of the KB requires some method of indexing information by agent (trying to minimize the amount of redundancy introduced) and should include rules for determining how *belief* spaces overlap, how to resolve conflicts of belief, etc. While there has been much work done on modal logics of knowledge and belief, few of the necessary primitives have yet been incorporated into extant KR schemes.

6.5.4. Time and States of Affairs

Standard applications assume that assertions are only about the "real" world, that is, a description or predicate is asserted to apply to a particular entity in some state of affairs. No other descriptions (e.g., general propositions, configurations and descriptions of states of affairs) can be asserted and no variable referents are allowed. However, the propositional operators of time and modality included in our MR require conditions other than those in the current state of affairs to be represented. Movement among states of affairs should be regulated by a set of accessibility relations, and propositions must be referred to as conceptual objects without asserting their (current) satisfiability.

Different states of affairs are not independent, disjoint worlds but are highly interrelated and often share the same set of definitional knowledge. They will vary according to which entities (mostly extensional) are believed to exist, which descriptions apply to such entities, and which propositions relate the entities. Compatible states of affairs then share the same taxonomy of conceptual frames, but will have constraints specifying their assertional differences.¹⁹ In terms of KL-One, this means having a fixed conceptual taxonomy with contexts, description wires, and nexuses organized to account for different states of affairs.

Time is a dimension which is given special status. It is a factor in all discourse, and has the most well-defined criterion for comparison, a partial ordering of times indices (or intervals). The task of developing a set of time modelling primitives in KR is a major research problem. Most approaches use ordering relations such as *before*, *after*, *during*, and *same-time* to structure their database. I have no concrete contributions in this area but refer the reader to relevant work. Allen (1981) constructs a hierarchical organization of intervals to maintain temporal relationships; McDermott (1982) has developed a temporal logic for reasoning about processes and planning; Mays (1985) provides a comprehensive survey of the area and a set of operators and axioms for temporal logics.

For other types of variability in states of affairs, there are less well-specified dimensions of change such as possibility, degree of obligation, levels of confidence and likelihood. It is not clear how to connect the different states of affairs along these dimensions. One promising approach is to permit states of affairs to be referable objects in the KR scheme and set up a taxonomy of conceptual frames to describe them. These "state of affair" frames encode the compatibility constraints that specify how the various contexts are related. Thus, they would include slots to define the conditions of a constraint, its modality, priority and degree of relevance in the current state of affairs, and be connected by links of similarity and difference. For certain dimensions such as levels of confidence, it is better to attach such information directly to the affected node in the KB. For

¹⁹ Certain contexts will warrant more radical changes in states of affairs for which common assumptions and definitions are suspended, for example, in representing fictional literature (e.g., *A Clockwork Orange*). These are

example, weighted roles could be attached to propositional nodes to indicate an agent's degree of conviction about its current assertional strength (Weiner 1984).

Many KR schemes include primitives for partitioning the knowledge base into "contexts" but few are sophisticated enough to model states of affairs. We have seen that the context mechanism of KL-One can be used to denote states of affairs but provides little means for structuring them (either internally or externally). Prolog/KR augments standard Prolog with a possible world mechanism which partitions hierarchically into "worlds"; it thereby allows dynamic, context-dependent inheritance and redefinition of predicate definitions. Primitives are included to specify the current active set of worlds and to specify changes of worlds. KRL defines *contingency* descriptors which specify a time or hypothetical world (or ranges of these) in which the descriptor is valid.

6.5.4.1. Choosing the Current State of Affairs

Independent of their organizational structure, the KB interpreter must decide what state of affairs is to be the current one. Fodor (1979) has suggested classifying predicates as *same-world* or *cross-world*. Same-world predicates require that all its arguments have referents in the same state of affairs (e.g., *kick*, *stand next to*), while cross-world predicates can relate referents in different states of affairs (e.g., *pray to*, *resemble*, *imagines*). Expanding on this, she identifies three principles for choosing the state of affairs in which a sentence should be evaluated. The *real-world* principle says that sentences and descriptions should be related to the real world whenever possible. Secondly, states of affairs in which a sentence has a truth value should be favoured. This is called *maximizing the truth value*. Finally the *world shifting* principle gives rules for determining the appropriate state of affairs for interpreting noun phrases. For instance, an indefinite noun phrase is taken to refer to an individual in the same world as the NP on its left, if there is one; otherwise an individual in the real world is preferred. These principles are only guidelines but nonetheless should be implemented as part of the reasoner/post-processor to control the preferred states of affairs.

6.6. Conclusions and Suggestions for Further Research

We have catalogued the possible ambiguities of opaque contexts, and defined a meaning representation to distinguish among the various readings. The traditional problems, namely the failure of substitution of equals and existential generalization, have been reformulated in a computational framework which requires evaluating descriptors and propositions along certain dimensions. Our work has emphasized the effects of opacity on noun phrases. For substitution of NP descriptors, the most important factors are the intensionality of the referent, the agents to whom the descriptor is ascribed, and the time of the descriptor application. In less detail, we also considered the effects of opacity on sentential constituents. When interpreting the propositional content of an embedded sentence, one must consider the time, strength, and modality of its assertion, the state of affairs it is relative to, and the agents to whom it is attributed.

The predicates and operators I have defined in the MR are intended to exemplify general cases and their semantics are specified at a correspondingly abstract level. No attempt has been made to ensure either that the set is complete, nor that all are necessary as primitives. For practical implementations, one would probably want to restrict the *types* of predicates and operators available, but provide a more complete set within each class and specify the semantics of individual predicates more succinctly (i.e., at a lower level). In terms of *formal* semantics, the MR could be given a more rigorous logical foundation by integrating it with current work on modal temporal logics and logics of knowledge and belief.

considered non-compatible states of affairs for which new concepts and new descriptive terms must be defined.

In any given opaque context, verbs and other constituents limit the possible readings that their arguments may have. For example, in Chapter 5 we discussed how copula verbs affect the interpretation of their subject and object noun phrases. This was of particular interest because copulas often establish equivalences among terms for subsequent substitution. However, there are many other intrasentential constraints among constituents (*e.g.*, same-world, same-time constraints (Fodor 1979)) that need to be investigated and added to restrict the set of possible readings further.

Given a set of candidate representations for an opaque context, there remains the problem of deciding which is the most correct. This will involve a number of pragmatic issues such as determining the intent and expectations of the speaker, maintaining a full discourse model, and selecting appropriate non-linguistic knowledge to apply in the discourse. For most opaque contexts, there is sufficient knowledge in the discourse environment to ensure a unique reading, but it is difficult for a NLP agent (*i.e.*, a non-human) to ascertain efficiently what parts of it are or are not relevant. Ambiguities of opacity reflect subtle distinctions in meaning and will require a sophisticated pragmatic component to distinguish among them.

We have considered what substitutions are *possible* in opaque contexts, but have not really considered what substitutions one would *want* to make in particular linguistic situations. Some reasons were mentioned in Chapter 5 for wanting to do substitution in an NLP system, namely to extract the maximum information from the input or to rephrase a clause for a third party. A NLP system should have a set of rules advising it when such substitutions would be preferred. For example, in Section 5.5.1 a rule was suggested for the substitution of labels and descriptors that depends on what the speaker knows and what he believes the listener knows. The descriptor predicates we have defined should prove useful for formulating these type of requirements.

Bibliography

References

- (Ali 1985).
Ali, Yawar, "Understanding Adjectives," Master's Thesis, University of Toronto, January 1985.
- (Allen 1981).
Allen, James, "An Interval-Based Representation of Temporal Knowledge," *Proceedings IJCAI-81*, vol. 1, pp. 221-226, Vancouver, B.C., August, 1981.
- (Allen 1979).
Allen, James F., "A Plan-based Approach to Speech Act Recognition," Report # 1131/79, University of Toronto, February, 1979.
- (Belnap 1975).
Belnap, N.D. and Anderson, A.R., *Entailment, The Logic of Relevance and Necessity*, Princeton University Press, 1975.
- (Bobrow 1977).
Bobrow, Daniel and Winograd, Terry, "An Overview of KRL, A Knowledge Representation Language," *Cognitive Science*, vol. 1, pp. 3-46, 1977.
- (Brachman 1983).
Brachman, Ronald, "What IS-A Is and Isn't: An Analysis of Taxonomic Links in Semantic Networks," *Computer*, vol. 16, no. 10, pp. 30-36, IEEE Computer Society, October 1983.
- (Brachman 1979).
Brachman, R.J., "On the Epistemological Status of Semantic Networks," in *Associative Networks: Representation and Use of Knowledge by Computers*, ed. N.V. Findler, pp. 3-50, Academic Press, New York, 1979.
- (Carnap 1947).
R. Carnap, *Meaning and Necessity*, University of Chicago Press, Chicago, 1947.
- (Church 1951).
Church, A., "A Formulation of the Logic of Sense and Denotation," in *Structure Method and Meaning: Essays in Honor of Henry M. Sheffer*, ed. P. Henle et al., Liberal Arts Press, New York, 1951.
- (Creary 1979).
Creary, Lewis, "Propositional Attitudes: Fregean Representation and Simulative Reasoning," *IJCAI-79*, 1979.
- (Dahl 1981).
Dahl, Veronica, "Translating Spanish into Logic Through Logic," *American Journal of Computational Linguistics*, vol. 7, no. 3, pp. 149-164, 1981.
- (Dowty 1981).
Dowty, David, Wall, Robert, and Peters, Stanley, *Introduction to Montague Semantics*, Dort Reidel Publishing, 1981.
- (Fodor 1981).
Fodor, Jerrold, "Representations," *M.I.T. Press*, 1981.
- (Fodor 1979).
Fodor, Janet Dean, *The Linguistic Description of Opaque Contexts (Outstanding Dissertations in Linguistics : 13)*, Garland Publishing, 1979. Originally author's PhD thesis.
- (Fodor 1980).
Fodor, Janet Dean, *Semantics: Theories of Meaning in Generative Grammar*, Harvard University Press, Cambridge, Mass., 1980. The Language and Thought Series.
- (Frege 1875).
Frege, Gottlieb, "On Sense and Reference," in *The Logic of Grammar*, ed. Donald Davidson and Gilbert Harman, pp. 116-128, Dickenson Publishing, 1975.

- (Friedman 1978).
Friedman, Joyce and Warren, David Scott, "A Parsing Method for Montague Semantics," *Linguistics and Philosophy*, vol. 2, pp. 347-372, 1978.
- (Grice 1975).
Grice, Paul, "Logic and Conversation," in *Syntax and Semantics (Vol. 3)*, ed. P. Cole and J. Morgan, pp. 41-58, Academic Press, New York, 1975.
- (Hirst 1986).
Hirst, Graeme, *Semantic Interpretation and the Resolution of Ambiguity*, Cambridge University Press, 1986.
- (Hofstadter 1982).
Hofstadter, Douglas R., Clossman, G. A., and Meredith, M. J., *Shakespeare's Plays Weren't written by him, but by someone else of the same name - An Essay on Intensionality and Frame-based Systems*, Indiana Linguistics Club, Bloomington, Indiana, November, 1982.
- (Jackendoff 1983).
Jackendoff, Ray, *Semantics and Cognition*, M.I.T. Press, 1983.
- (Kaplan 1978).
Kaplan, David, *On the Logic of Demonstratives*, pp. 401-412, 1978.
- (Karttunen 1977).
Karttunen, Lauri and Peters, Stanley, "Conventional Implicature," in *Syntax and Semantics (Vol. 11) - Presupposition*, ed. Choon-Kyer Oh and David Dineen, 1977.
- (Kripke 1963).
Kripke, S., "Semantical Considerations of Modal Logic," *Zeitschrift fur Mathematische Logik and Grundlagen der Mathematik*, vol. 9, pp. 67-96, 1963.
- (Levesque 1979).
Levesque, Hector and Mylopoulos, John, "A Procedural Semantics for Semantic Networks," in *Associative Networks : Representation and Use of Knowledge by Computers*, ed. N.V. Findler, pp. 93-126, Academic Press, New York, 1979.
- (Levesque 1983).
Levesque, Hector, "A Logic of Implicit and Explicit Belief," *AAAI-83*, pp. 198-202, 1983.
- (Levinson 1983).
Levinson, Stephen C., *Pragmatics*, Cambridge University Press, Great Britain, 1983.
- (Lipkis 1982).
Lipkis, Tom, "A KL-One Classifier," *Proceedings of 1981 KL-One Workshop*, pp. 128-145, Bolt Beranek and Newman, June 1982.
- (Maida 1982).
Maida, Anthony and Shapiro, Stuart, "Intensional Concepts in Propositional Networks," *Cognitive Science*, vol. 6, pp. 291-330, 1982.
- (Mallery 1985).
Mallery, John C., "Universality and Individuality : The Interaction of Noun Phrase Determiners in Copular Clauses," *Proceedings 23rd ACL Meeting*, pp. 35-41, Chicago, May 1985.
- (Mark 1982).
Mark, William, "Realization," *Proceedings of 1981 KL-One Workshop*, pp. 78-89, Bolt Beranek and Newman, June 1982.
- (Mays 1985).
Mays, Eric, "Some Conditions on Providing a Class of Extended Responses," *Proceedings Theoretical Approaches to Natural Language Understanding*, CSCSIO, Halifax, Nova Scotia, May 1985.
- (McDermott 1982).
McDermott, Drew, "A Temporal Logic for Reasoning About Processes and Plans," *Cognitive Science*, vol. 6, 1982.

(Montague 1973).

Montague, Richard, "The Proper Treatment of Quantification in Ordinary English," in *Approaches to Natural Language: Proceedings of the 1970 Stanford Workshop on Grammar and Semantics*, ed. Hintikka, Kaarlo Jaakko Juhani; Moravcsik, Julius Matthew Emil and Suppes, Patrick Colonel, pp. 221-242, Dort Reidel, Dordrecht, 1973. Also in *Formal Philosophy: Selected Papers of Richard Montague*. Richard Hunt Thomason, New Haven: Yale University Press, 1974. pp. 247-270.

(Moore 1981).

Moore, Robert, "Problems in Logical Form," *Proc. 19th ACL Conference*, pp. 117-124, Stanford, 1981.

(Nakashima 1984).

Nakashima, Hideyuki, "Knowledge Representation in Prolog/KR," *Symposium on Logic Programming*, pp. 126-130, IEEE, 1984.

(Nishida 1983).

Nishida, Toyooki, *Studies on the Application of Formal Semantics to English-Japanese Machine Translation*, Department of Information Science, Kyoto University, October 1983. Doctoral Thesis.

(Pereira 1984).

Pereira, Fernando, "C-Prolog User's Manual," Technical Report, SRI International, Menlo Park, California, 1984.

(Popowich 1984).

Popowich, Fred, "SAUMER: Sentence Analysis Using MEtarules," LCCR TR 84-2, CMPT TR 84-10, Simon Fraser University, Burnaby, B.C., August 1984.

(Popowich 1985).

Popowich, Fred, "The SAUMER User's Manual," LCCR TR 85-4, Laboratory for Computer and Communications Research, Simon Fraser University, Burnaby, B.C., March 1985.

(Quirk 1984).

Quirk, Randolph and Greenbaum, Sidney, *A University Grammar of English*, Commonwealth Printing Press Ltd., Hong Kong, 1984.

(Schmolze and Brachman 1982).

Schmolze, J.G. and Brachman, R.J., "Proceedings of the 1981 KL-One Workshop," Report No. 4842, Bolt Beranek and Newman Inc., June 1982.

(Schubert and Pelletier 1982).

Schubert, Lenhart and Pelletier, Francis, "From English to Logic: Context-Free Computation of Conventional Logical Translation," *American Journal of Computational Linguistics*, vol. 8, no. 1, pp. 27-44, Jan-March 1982.

(Stalnaker 1972).

Stalnaker, Robert, "Pragmatics," in *Semantics of Natural Language*, ed. Donald Davidson and Harman, pp. 380-397, 1972.

(Strzalkowski and Cercone 1985).

Strzalkowski, Tomek and Cercone, Nick, "A Framework for Computing Extra-Sentential References," *Proceedings Theoretical Approaches to Natural Language Understanding*, CSCSI/SCEIO, Halifax, Nova Scotia, May 1985.

(van Benthem 1985).

Van Benthem, Johan, *A Manual of Intensional Logic*, Center for the Study of Language and Information, Stanford University, California, 1985.

(Warren 1982).

Warren, D. H. D., "Higher Order Extensions to PROLOG - Are They Needed?," *Machine Intelligence*, vol. 10, 1982.

(Warren 1983).

Warren, David Scott, "Using Lambda Calculus to Represent Meanings in Logic Grammars," Report #83/044, SUNY at Stony Brook, January 1983. Also in Proceedings, 21st Annual Meeting, Association for Computational Linguistics, Cambridge MA, June(1983).

(Weiner 1984).

Weiner, E. Judith, "A Knowledge Representation Approach to Metaphors," *Computational Linguistics*, vol. 10, no. 1, 1984.

(Winograd 1976).

Winograd, Terry, "Towards a Procedural Understanding of Semantics," Report Number STAN-CS-76-580, Stanford University, November 1976.

(Woods 1975).

Woods, William, "What's in a Link: Foundations for Semantic Networks," in *Representation and Understanding: Studies in Cognitive Science*, ed. D.G. Bobrow and A. Collins, pp. 35-82, Academic Press, New York, 1975.

APPENDIX A

Classification of Noun Phrases

1. Definite Descriptions

The following types of noun phrases are considered as definite descriptions.

- (i) Proper Nouns:
Proper nouns label particular people, places, publications, and the like. They are always referential.¹
Ex. *John Lennon, King George V, Trafalgar Square*
- (ii) Pronouns:
Noun phrases consisting only of a personal pronoun (*he, she, it,...*), demonstrative pronoun (*that, those, this, these*), or possessive pronoun (*his, hers, mine,...*) are always referential. They indicate some particular entity that the speaker and hearer are mutually aware of.
- (iii) Common Nouns preceded by a Demonstrative Article:
Noun phrases having a demonstrative article (*this, that, these, those*) preceding a common noun are similar to type (ii) pronouns. They point out a particular entity and can be only read referentially.
Ex. *That red boat, These children, That ice-cold water*
- (iv) Common Nouns preceded by a Definite Article or Possessive Article:
Any noun phrase which includes either the definite article *the* or any of the possessive articles (*my, yours, his,...*) is considered a definite description. These are often ambiguous and can be read either referentially or attributively. Certain quantified noun phrases are included in this group, in particular those containing a definite article from the list above.
Ex. *his favourite aunt, all my children, the furniture I bought*

2. Indefinite Descriptions

Indefinite descriptions fall into one of the following categories.

- (i) Pronouns:
Question (*who, what,...*) and quantified (*someone, something*) pronouns are considered indefinite and always read non-specifically.
Ex. *Who came to dinner?*
Someone will phone you today.
- (ii) Common Nouns preceded by a Indefinite Article:
Using *a(n)* followed by a common noun creates an indefinite noun phrase that is ambiguous with respect to specificity. Similar comments apply to question NPs consisting of

¹ In logic these are often termed rigid designators since they are interpreted as denoting the same individual across all possible indices.

which or *what* followed by a common noun.

Ex. *A young child* walked into the room.

Which boy ate my lunch?

(iii) Mass Nouns and Generic Noun Phrases:

Mass nouns refer to an undifferentiated mass or continuum (*e.g.*, water). They have no plural form and often are used without determiner². *Generic* nouns, which syntactically take the form of unmodified plural count nouns or singular count nouns preceded by an article, are used to refer to a class and so denote what would be normal or typical for members of that class.

Ex. *The cake* was dry and crumbly. (Definite Noun)

Cake is a wonderful source of energy. (Mass Noun)

A lion is a fierce fighter. (Generic Noun)

A lion sat on the hood of my car. (Count Noun)

Since these indefinite descriptions are used to refer to a *class* of entities rather than to the individual members belonging to that class, both generic and mass nouns are interpreted as non-specific.

(iv) Indefinite Quantified Noun Phrases:

Noun phrases modified by quantifiers (*every, each, all, some, a few, several, one, two, many,...*) which do not include a definite determiner will be considered as indefinite descriptions.³ These are ambiguous with respect to specificity.

Ex. *All children, Many people, A few grapes*

² In contrast, *count* nouns represent individual countable entities. They can be singular or plural but the singular must be preceded by some type of determiner. Some nouns such as "cake" and "paper" can belong to either category.

³ There is another dimension of ambiguity which arises in this category and with plural definite descriptions (*e.g.*, the two boys), namely whether the noun phrase is to be read distributively or collectively.