1 Problems in the semantic analysis of text

1.1 Introduction

A primary problem in the area of natural language processing is the problem of semantic analysis. This involves both formalizing the general and domain-dependent semantic information relevant to the task involved, and developing a uniform method for access to that information. Natural language interfaces are generally also required to have access to the syntactic analysis of a sentence as well as knowledge of the prior discourse to produce a detailed semantic representation adequate for the task.

Previous approaches to semantic analysis, specifically those which can be described as using *templates*, use several levels of representation to go from the syntactic parse level to the desired semantic representation. The different levels are largely motivated by the need to preserve context-sensitive constraints on the mappings of syntactic constituents to verb arguments. An alternative to the template approach, *inference-driven mapping*, is presented here, which goes directly from the syntactic parse to a detailed semantic representation without requiring the same intermediate levels of representation. This is accomplished by defining a grammar for the set of mappings represented by the templates. The grammar rules can be applied to generate, for a given syntactic parse, just that set of mappings that corresponds to the template for the parse. This avoids the necessity of having to represent all possible templates explicitly. The context-sensitive constraints on mappings to verb arguments that templates preserved are now preserved by filters on the application of the grammar rules. This allows a more concise and extendable representation of the verb semantics for a given domain since advantage can be taken of linguistic generalizations about certain syntactic mappings.

Two other important characteristics of inference-driven mapping will also be presented here. The use of one level of representation allows the interleaving of several different semantic and pragmatic subtasks. The decompositional *semantic predicate* level of representation will also be described in detail, as well as the way that it makes explicit the verb-independent effect that certain semantic roles, such as INSTRUMENTS, can have on the final representation of a sentence.

Inference-driven mapping is specifically designed for finite, well-defined, i.e.,

limited, domains about a particular topic. The coverage of this type of domain is clearly delineated, and the processes required for the application are straightforward enough to be formalizable, making them especially appropriate for the use of the closed-world assumption [Pereira and Shieber, 1987]. Under the closed-world assumption, a specific application domain is formalized with axioms that allow the truth of a predicate to be proven. If a predicate cannot be proven to be true, it is assumed to be false; i.e., negation as failure [Clark, 1978]. Information that is not relevant to the domain does not need to be formalized, thus allowing the implementer to sidestep the insurmountable task of representing large amounts of general world knowledge.

An example of such a domain is a set of physics word problems for college students involving pulley systems, the *pulley domain*. Each problem is stated in English sentences that completely describe a miniature world of physical entities and relationships between those entities. The goal of the natural language processor is to produce a semantic representation of each problem that is sufficiently detailed to enable a computer program to produce the correct solution of the problem [Bundy, 1979]. This final semantic representation consists of a set of partially instantiated logical terms known as *semantic predicates*.

The formalization of the domain is essential for solving the following basic tasks which are associated with the semantic processing of text:

1. establishing referents for the noun phrases;
2. finding appropriate mappings from the syntactic constituents of the parse onto the underlying semantic representation of the verb;
3. using pragmatic information to assign fillers to semantic roles that do not have an explicit syntactic realization (the term "pragmatic" is used to refer to both discourse knowledge and general and domain-dependent information);
4. expanding the representation of the verb into a more detailed representation that fulfills the requirements of the processing task;
5. constraining allowable inferences so that this semantic representation does not become unmanageably large;
6. appropriately integrating the final representation of the clause with the representations of prior clauses.

These tasks will all be described in more detail, but first certain elementary terms must be defined.

1.2 The semantic representation of sentences

Deriving an appropriate semantic representation for a single sentence in a given context is a difficult problem in natural language processing. It is a non-trivial task, even in a limited domain where one is restricted to discussing

inanimate entities and can assume the simplest case of meaning, i.e., where it is assumed that a statement about inanimate entities has a strict meaning and the speaker intends it to be interpreted literally. Given a syntactic parse and a consensus on what the semantic representation should be, there is still the problem of assigning the correct semantic role to each syntactic constituent of the parse, and producing the indicated semantic representation. This requires:

1. formalization of domain-specific information;
2. knowledge of the different syntactic cues that can be used to indicate semantic roles;
3. pragmatic information about the entities mentioned.

This section explores the difficulties associated with each of these requirements, illustrated with specific examples from the pulley domain. Before discussing these factors in detail, examples of the use of **syntactic constituent** and **semantic role** are given.

1.2.1 Examples using syntactic constituent and semantic role

The following discussion summarizes one of the more popular linguistic analyses of the use of "subject." It covers three recognized uses of subject: grammatical (G), logical (L), and thematic (T). They can be distinguished by the following examples:

1. John took the largest dog.
   John is assigned G, L, T

2. The largest dog was taken by John.
   dog is assigned G and T
   John is assigned L

3. The largest dog, John gave away.
   dog is assigned T
   John is assigned G and L

The grammatical subject usually immediately precedes the main verb, agrees with it in number, and is not marked by a preposition. The logical subject typically corresponds to the first argument of the logical predicate chosen to represent the verb. The thematic subject indicates the "theme" of the sentence, and can be said to correspond to the topic or focus. It is usually present as the first noun phrase. Generally, "subject" is used here to mean the grammatical subject, unless indicated otherwise. The syntactic category "subject" along with the referent of the noun phrase that the category is associated with, such as John, are indicated by subj(john), where subj is a function symbol with one argument, and john is the instantiation of that argument.

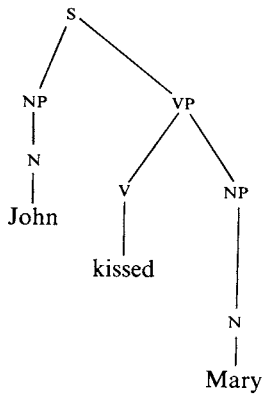

Fig. 1.1. Syntactic parse

The syntactic structure of "John kissed Mary" is given in figure 1.1.

Generating the parse tree from the string "John kissed Mary" is known as syntactic processing or syntactic parsing, and the noun phrases designated as subject and object are syntactic constituents. In a normal, active, declarative sentence in English, given an unambiguously transitive verb such as ***kiss***, the semantic role AGENT can be assigned to "John" and the semantic role PATIENT to "Mary." This is intended to capture the notion that "John" is doing the "kissing" and "Mary" is the one being "kissed." A general semantic representation of ***kiss*** could be ***kiss***(agent,patient). The correspondence between the general representation and the representation of the particular parse can be captured with the following rules:

Replace agent with john if subj(john)
Replace patient with mary if obj(mary)

This results in the following semantic representation for the sentence:

kiss(john,mary).

"John kissed Mary" is clearly quite different from "Mary kissed John." The ordering of the words reflects a particular syntactic structure that strongly affects what is being communicated. In "Mary kissed John," the semantic roles assigned to the syntactic constituents would be

Replace agent with mary if subj(mary)
Replace patient with john if obj(john)

resulting in the following semantic representation:

kiss(mary,john).

1.2.2 The inherent difficulties in defining semantic representations

The following sections describe two different types of constraints on the formation of logical predicates for producing semantic representations: (1) the demands of the task to be performed for which the representation is required, and (2) the requirements of the process that does reference evaluation and fills semantic roles. For a semantic processor these constraints are actually of great benefit, since the consensus on the nature of semantic representation referred to so lightly above is not so readily obtainable. The next two sections introduce two of the major problems associated with defining semantic representations: (1) the inherent indivisibility of words and (2) verb ambiguities.

Defining primitives

There are essentially two different linguistic approaches to providing semantic representations: semantic markers and decomposition [Jackendoff, 1970]. Both these methods are based on the recognition of fundamental similarities between the ways in which many words can be used. A classic example of the use of semantic markers is the definition of "bachelor" by the markers UNMARRIED and MALE. The standard example for decomposition is the decomposition of "kill" into CAUSE to DIE. The motivation for both of these approaches is to define complex terms as collections of simpler terms, which may themselves be decomposed. Eventually a set of primitive terms will be reached, but this does not end the process. It is not clear that the primitive terms, whatever they may be, can or should be defined [Fodor, 1980]. The difficulty, indeed impossibility, of providing a semantic analysis of language is a well-worn philosophical enigma.

The "vagueness' of semantic notions can easily be demonstrated using the example of "John likes Mary." Given that a "like" relationship exists between "John" and "Mary," i.e. ***like***(john,mary), what inferences should then be drawn? Defining the relationship more precisely is very difficult, especially for a computer that has no common fund of human emotional experience to draw upon. Using dictionary definitions such as "fondness," "being pleased with" "having a preference for" merely defers the issue.

The question still remains: Is it possible to derive appropriate semantic representations for verbs, nouns and prepositions that can be composed to form appropriate semantic representations for sentences? The key to answering the question is to sidestep the issue of the compositionality of all of English, and limit the problem to providing a computer with tools that are adequate for a given task. The issue then becomes not "What is the meaning of X?", but rather "What do we need to know about X's meaning in order to do Y?" What is assumed about X should be a proper part of X's "meaning," but should not in any way be expected to give predictions about <u>all</u> of X's meaning.

This is the approach to providing semantic representations that was first

exemplified in the pathbreaking implementations of Winograd [1972] and Woods [1973], as described in chapter 2. It has been the cornerstone for much of the research in knowledge representation that has followed [Halpern, 1986; Brachman and Levesque, 1985].

Limited domains

Establishing a useful subpart of X's meaning is a complicated problem in itself. It is somewhat simplified by strictly limiting the context in which X is to be used, so that multiple meanings can be avoided, as well as metaphorical uses. The idealized pulley world described in section 1.3 is an example of a simple, finite, concrete domain. Entities and the relationships between them are clearly defined. In such a domain, a *limited domain*, axiomatizing the information about X that is needed in order to perform Y is a more tractable task.

Verb ambiguities

Another major problem in determining appropriate semantic representations is caused by verb ambiguities; the same verb having more than one definition. The semantic features of the subject or the object of the verb can usually determine how the verb is being used, as in the examples given below. Every new object for ***throw*** results in a different interpretation. To complicate matters further, "throw a boxing match" could conceivably have two interpretations, "organize a boxing match," and "purposely lose a boxing match"; an ambiguity that would have to be resolved by context. Clearly, the predicates chosen to represent ***throw*** or ***take*** would have to change radically depending on the nature of the object. This phenomenon is not restricted to objects. Three different senses of ***run*** are unambiguously indicated by the subject in the last three examples.

throw a baseball
throw support behind a candidate
throw a boxing match
throw a party
throw a fit

take a book from the shelf
take a bus to New York
take a nap
take an aspirin for a cold
take a letter in shorthand

the clock runs
John runs
my nose runs

If the semantic features associated with expected semantic roles are used to choose between alternative verb definitions, then the selection process does not necessarily precede the mapping process, but may have to proceed in parallel with it.

In the pulley domain, human emotions and intentionality are ignored and ambiguous verbs such as ***throw*** and ***take*** can be avoided. Yet, the task is by no means trivialized. The problem of verb ambiguities is restricted to determining whether similar uses of the same verb represent separate definitions or the same definition. For example, do the uses of ***hang*** and ***suspend*** in the following clauses constitute several different definitions of the verbs or are they all really the same?

a pulley is suspended from a pulley
at one end of a string a weight is suspended
a string hangs over a pulley
weights hang freely on the ends of a string

In changing from the causative form of ***hang***, which includes the optional AGENT role, to the stative form which does not, is there a different definition of the verb involved?

John hung a mass from a pulley.
A mass hangs from a pulley.

Chapter 5 discusses in detail how the system presented here handles each of these verb usages, and suggests that while for inference-driven mapping they each involve the same definition of the verb, the semantic representation changes radically from phrase to phrase because of the effect of the fillers of the semantic roles.

1.2.3 Mapping between syntactic constituents and semantic roles

A major task of semantic analysis is to provide an appropriate mapping between the syntactic constituents of a parsed clause and the semantic roles associated with the verb. Three factors complicate the mapping:

1. the ability of the syntactic constituents to indicate several different types of semantic roles given appropriate contexts;
2. the large number of choices available for syntactic expression of any particular semantic role;
3. preposition ambiguities.

The factors preclude the possibility of a one-to-one mapping between the syntactic constituents and the semantic roles. The pulley domain contains examples of many of these complications but not all of them. It is important to present them in their entirety here, however, since they provided important

motivation for the design of previous semantic processors, and since inference-driven mapping is designed to handle them. Problems that are not illustrated fully by the pulley domain, such as semantic-role interdependencies, are discussed in terms of the examples given here, and the techniques by which inference-driven mapping deals with them are described in chapter 3.

Optional and obligatory semantic roles

The first step in the mapping process is the selection of the relevant verb definition. These definitions consist mainly of the appropriate semantic predicate and its associated semantic roles. All of the semantic roles may not be filled in a particular usage of a verb, but this does not necessarily imply a new verb definition. Several verbs such as ***open*** and ***break*** have **optional** semantic roles, which means that the roles may or may not appear in the surface structure of a clause containing the verb. Given {AGENT, INSTRUMENT, and PATIENT} as the set of expected roles for ***open***, the phrases listed below illustrate the **optional** occurrence of the AGENT and the INSTRUMENT roles:

Optional semantic roles

John opened the door with a key.
The door was opened by John.
The door was opened with a key.
A key opened the door.
The door opened.
John gave Mary the book.
John gave the book to Mary.

The same sets of roles can be expressed using different constituents, as in the last two examples using ***give***. The semantic roles seem to be playing musical chairs with the available syntactic constituents, with one or more roles often being omitted. For ***open***, the only **obligatory** role is the PATIENT role that is filled by the door. For ***give***, all three roles are usually obligatory. *In summary, more than one syntactic constituent may indicate a particular role, and conversely more than one role may be indicated by a particular constituent.*

The associations between semantic roles and syntactic constituents, although complex, are not arbitrary. There is general agreement that for English many verbs use at least three semantic roles, AGENT, PATIENT and INSTRUMENT, and there are indications of rules of etiquette to be followed in their possible mappings to syntactic constituents. These rules are exemplified by the *precedence relations* expounded by Fillmore [1977] in his theory of case. In an active sentence, if the AGENT is present it is the subject, or else the INSTRUMENT (if present) is the subject, or else the PATIENT is the subject. This evidence of some regularity in semantic role assignment is of paramount importance to anyone attempting to understand the nature of the relationship

between syntax and semantics. It is essential information for a processor that performs mappings from semantic roles to syntactic constituents.

Some of the examples of alternative syntactic realizations of verbs in the pulley domain are given below. The semantic processor accepts all the possible expressions and produces appropriate representations for them.

Alternative syntactic realizations

A particle is connected to another particle by a string.
A string connects two particles.
A particle is attached at the end of string.
A particle is attached to the end of a string.
A string with particles attached . . . .

Preposition ambiguities

Having seen the subject observe a certain degree of decorum with respect to role assignment, it is tempting to look for other useful regularities in the performance of prepositional phrases. It is true that an INSTRUMENT can be expected to be indicated by either the subject of a BY or a WITH prepositional phrase. But WITH and BY are by no means restricted to introducing instruments. BY can also indicate an AGENT, and WITH is often used as a **comitative**, indicating someone that goes along with someone else, or as a **locative**, indicating a location. This is the case even for verbs that expect INSTRUMENTS. A verb's possible semantic roles do not restrict the use of prepositions in association with that verb. Illustrations of the variety of uses of WITH that can appear with ***open*** and other INSTRUMENT-taking verbs are given below.

Prepositions performing independently of verb expectations

comitative:
John opened the door with Mary and Jim.

instrumental:
The door was opened with a key.

manner:
The door opened with a solid click.

instrumental:
The door opened with a solid whack.

manner:
John hit the door with a solid whack.

manner:
John kicked the door with a solid whack.

However, the way a preposition can be used is not entirely independent of the verb. It might seem plausible that semantic features on the object of the preposition could exclusively determine its use, but this is clearly not the case. The dividing lines between comitative uses, locative uses, and instrumental uses are hard to draw, but wherever they are, they are strongly affected by the verb semantics. The sentences below give examples of WITH prepositional phrases that can have the same type of object and yet still have a variety of uses depending on the verb they appear with.

Preposition use being determined independently of semantic features on the object

comitative:
Mary started the introductions with John.

instrumental:
The townspeople filled the gap in the firebucket line with the town drunk.

instrumental (?):
Mary flirted with John to get even with Bill.

locative:
Mary put the books with the papers.

instrumental:
Mary stuffed the hole in the window with old newspapers.

comitative:
A kite with paper streamers floated into view.

Semantic role interdependencies

A legacy of the precedence relations mentioned earlier with respect to semantic roles is the constraints they place on prepositions. In spite of the ability of both BY and WITH to introduce the INSTRUMENT role, they cannot typically be substituted for one another. This is illustrated by the following examples.

John opened the door {with/*by} a key.
The door was opened {by/*with} John.

There is a partial explanation for this in the subtle ways in which semantic roles defer to other semantic roles in the assignment of syntactic constituents. The AGENT takes precedence over the INSTRUMENT where the subject is concerned. It is also true that if the AGENT is the subject, the INSTRUMENT cannot usually be indicated by a BY but rather by a WITH. Neither can the INSTRUMENT