

Natural Language Processing & Applications

Meaning

1 Introduction

Processing meaning in NL is, not surprisingly, difficult. What do we mean by claiming that a person understands the meaning of a piece of NL? (Or to put another way, what do we mean by ‘meaning’?)

One answer is to use behavioural criteria.¹

1. If the NL involves an instruction, then the person to whom it was addressed understands it if he or she correctly carries out the instruction.
2. If the NL is informative, a person understands it if he or she is able to paraphrase the information, or use it to reach new conclusions.

However, there seem to be some further, possibly ‘mentalistic’ criteria.

- A dog can be trained to carry out quite complex commands (e.g. when rounding up sheep). Does the dog understand the commands? The unix operating system responds to commands. Does it understand them?
- A Prolog system can deduce new conclusions from a given piece of code. Does it understand the code?

Here we have strayed into some fascinating, but difficult and highly controversial areas of philosophy. As far as computer processing of the meaning of NL is concerned, it seems to me that only behavioural criteria are available. Hence I shall say that a system is processing the meaning of NL if it meets either of the behavioural criteria given above.

Earlier I divided ‘meaning’ into semantics and pragmatics. Semantics was essentially defined as ‘compositional meaning’, i.e. the meaning of a sentence based on the meaning of the words from which it are composed, plus the meaning associated with the ordering of words within the sentences. Hence by semantics, I mean, roughly, ‘literal’ meaning. Pragmatics, on the other hand, was defined as including the intentions of the NL user. Initially, I will consider only semantics.

2 Referential Semantics

One of the earliest systems to demonstrate successful syntactic and semantic processing of NL was Winograd’s (1972) SHRDLU. The basic idea is that we have a world in which there are blocks of different colours and shapes. SHRDLU only SIMULATED vision and movement, but we can imagine a robot having been programmed to identify colours and shapes and associate them with words. Thus when a block is presented to the robot’s vision system, it can identify its colour as *blue* or its shape as *pyramidal*. The robot is also programmed to be able to carry out specific movements, e.g. to move its arm to a particular location, to grasp, to lift, etc. It does this ultimately by the execution of machine code, but is equipped with a compiler or interpreter for a special ‘Robot Command Language’ (RCL). We give the robot the NL instruction *Put the small red pyramid on the large green cube*.

- Syntactic and semantic processing can be used to transform the NL instruction into an RCL instruction.
- The RCL instruction can then be executed.

The key point here is that words like *red* or *pyramid* have a meaning because they can be mapped onto the outputs of the robot’s built-in identification procedures. *Put* has a meaning because it can be mapped onto a sequence of the robot’s built-in movements. (The mapping need not necessarily be 1:1.)

Note that many of the problems involved in processing NL can be demonstrated even in such a simple system; e.g. the need for parallel syntactic and semantic processing. A command

¹ This is, of course, the essence of the Turing test.

such as *Put the red pyramid on the green cube on the yellow cube*, could mean either *Put (the red pyramid) on (the green cube (on the yellow cube))* or *Put (the red pyramid (on the green cube)) on (the yellow cube)*. The syntactic ambiguity can only be resolved by semantic processing. For example, if the world contains a green cube on a yellow cube but not a red pyramid on a green cube, then the first interpretation is correct.

The underlying approach seems to be:

1. Create a set of ‘primitives’ consisting of the identifications, movements, etc. of which the robot is capable.
2. Translate the NL input into these primitives. The ‘meaning’ of the NL then just IS the underlying primitive.

This approach has been called ‘referential semantics’ because the meaning of words is what they refer to. Thus the meaning of a noun phrase like *the green box* is one particular green box; the meaning of an adjective like *green* is the appropriate output of a colour identification program; the meaning of a verb like *put* is a set of actions that a robot performs.

Another example of this approach might be the interface to a database. The database contains information in the form of records whose fields describe entities. A NL query is translated into the appropriate database query language. Imagine a personnel database, and a query like *How much does Joan Smith earn?* The meaning of *Joan Smith* is the entity as modelled in the database; the meaning of *How much does ... earn?* is a query on the appropriate field/record.

Clearly this approach can account for only a tiny proportion of human uses of NL. How many of the sentences in this handout, for example, can be understood by referential semantics? However, it may be that a considerable fraction of the interactions we want to have with machines do involve situations where referential semantics is adequate.

Although semantic processing is incorporated into some practical NLP systems, it is much less obviously successful than speech and grammar processing.

- Lexical semantics is concerned with ways of representing the meaning of individual lexemes. A particular issue is determining which lexeme is represented by words which are written and/or spoken identically (homographs and homophones, collectively homonyms).²
- The dominant paradigm in representing meaning has been the use of logic, particularly first-order predicate calculus. This approach is introduced in the second year module, Natural Language Processing 1, and will not be covered again here: however you should be aware of its considerable importance.

In this part of the module I have chosen two relatively small parts of the topic in order to try to illustrate the problems (and successes) involved in processing meaning.

3 Semantic Features

In the approach described above, we began by defining a set of ‘primitives’ upon which meaning was then based. In referential semantics, these primitives can be understood as direct references to objects or to classifications or actions performed by some system. The next step seems to be to free the primitives from the referential restriction. How, for example, can we handle a noun like *boy*, which refers not to one specific object, but to a set of objects?

Nouns can be classified according to the values of a set of **semantic features**. For example, the word *boy* can be paraphrased as *young male human*, *man* as *not-young male human*. The difference in meaning between *boy* and *man* seems to reside in the value of age. Thus we could take AGE to be a primitive, here taking the values *young* and *not-young*. *Boy* is defined by AGE = *young* (among other features), *man* by AGE = *not-young*. A NL user (human or

² The word *bank* with meanings ‘river bank’ and ‘source of money’ is an example of a homonym which is both a homograph and a homophone. *Lead* as in *Kings no longer lead their men into battle* and *Lead is a dense metal* is an example of a homograph but not a homophone, whereas in *Kings led their men into battle* and *Lead is a dense metal*, [led] is a homophone but not a homograph.

computer) is assumed to be able to determine the values of primitive features INDEPENDENTLY OF THE LANGUAGE SYSTEM. The meaning of *boy* is then something like ‘a member of the set of objects which have the primitive features AGE = young, ...’.

Features are often binary. In this case, there is a convention that instead of saying that FEATURE has one of two values, we use +FEATURE as a shorthand for FEATURE = +, -FEATURE as a shorthand for FEATURE = -. Thus we could use +YOUNG and -YOUNG instead of AGE = young and AGE = not-young. (In the case of a noun for which YOUNG is not relevant, an alternative to -YOUNG is to use ±YOUNG or 0YOUNG, i.e. treat YOUNG as a THREE-valued feature.)

YOUNG is a rather specific feature. Some more general semantic features which have been used for nouns include:

ABSTRACT	Nouns with the feature +ABSTRACT are abstract or non-concrete (e.g. <i>sincerity</i>), those with the feature -ABSTRACT are concrete (e.g. <i>Jane</i> or <i>water</i>).
COMMON	All ‘proper nouns’ (names of people, things, etc.) are -COMMON. Thus <i>Jane</i> (referring to a specific person) is -COMMON, as is <i>London</i> . Other nouns are +COMMON, e.g. <i>dog</i> , <i>mankind</i> or <i>sincerity</i> .
COUNT	Nouns which can be made plural are +COUNT, e.g. <i>dog</i> . Nouns which are -COUNT include <i>water</i> (in its usual sense), <i>mankind</i> or <i>sincerity</i> .
ANIMATE	Nouns with the feature +ANIMATE are alive, those with -ANIMATE are not.
HUMAN	+HUMAN implies human, -HUMAN implies not human.
MALE	+MALE implies male, -MALE implies not male (i.e. either female or neither).
FEMALE	+FEMALE implies female, -FEMALE implies not female (i.e. either male or neither).

It’s arguable whether some of these features are ‘semantic’ or not; in this module I won’t distinguish between obviously semantic features such as HUMAN and those which are more syntactical such as COUNT.

Nouns can then be classified using sets of these primitive features. For example:

<i>Jane</i>	[-ABSTRACT, -COMMON, -COUNT, +ANIMATE, +HUMAN, -MALE, +FEMALE]
<i>boy</i>	[-ABSTRACT, +COMMON, +COUNT, +ANIMATE, +HUMAN, +MALE, -FEMALE]
<i>idea</i>	[+ABSTRACT, +COMMON, +COUNT, -ANIMATE, -HUMAN, -MALE, -FEMALE]
<i>sincerity</i>	[+ABSTRACT, +COMMON, -COUNT, -ANIMATE, -HUMAN, -MALE, -FEMALE]

These features are not independent, so that rules can be used to predict features from one another, e.g.

+ABSTRACT ⇒ [+COMMON, -ANIMATE] ⇒ [+COMMON, -HUMAN, -MALE, -FEMALE]
 -ANIMATE ⇒ [-HUMAN, -MALE, -FEMALE]
 +HUMAN ⇒ [+ANIMATE, -ABSTRACT]

Using such rules, descriptions can be simplified. It is adequate to describe *sincerity* as [+ABSTRACT, -COUNT].

Although I won’t discuss coding details here, it should be fairly obvious that we can store features in a lexicon, preferably the same one we use for syntax analysis, and then use ‘feature matching’ to restrict or extend the analysis. For example, in syntactic processing, we used the idea that verbs select certain SYNTACTIC categories as their complements. Thus:

**She put the bottle.*

was syntactically invalid because *put* selects the verb complements NP+PP. In a similar way, the use of semantic features can lead to the rejection of sentences such as:

**I saw his sincerities.*

**My brother is a married bachelor.*

**Ideas sleep.*

**The stones cried.*

Sincerity has the semantic feature -COUNT so cannot be made plural, so that *sincerities* will be rejected during morphological analysis. *Married* implies -SINGLE whereas *bachelor* implies +SINGLE, so they are contradictory. We can require the features of adjectives to match those

of the nouns or SNPs they qualify. The verbs *sleep* and *cry* require a subject with the semantic feature +ANIMATE, which can be checked during parsing.

Adjectives both possess their own properties and select for properties in the SNPs they modify, and the lexicon will need to reflect this fact. For example, *married* has the feature -SINGLE, and selects for +HUMAN, so that *married horse* is semantically odd. As another example, consider the sentences:

Although green, the girl's skirt was beautiful.
Although green, the apple's taste was beautiful.

Based on syntax alone, both sentences are ambiguous: the adjective *green* could refer to the either of the two nouns which follow. Using semantic features, we can remove the ambiguity. *Green* selects for [-ABSTRACT, -HUMAN]. *Girl* has the features [-ABSTRACT, +HUMAN], *skirt* the features [-ABSTRACT, -HUMAN]. Hence *green* must modify *skirt* not *girl*. *Apple* has the features [-ABSTRACT, -HUMAN], *taste* is [+ABSTRACT, -HUMAN]. Hence *green* cannot modify *taste* but can modify *apple*. It will be tedious to construct the lexicon (although this can be aided by rules for deducing features from one another), but once done, processing is basically a matter of pattern-matching.

However, this approach soon runs into problems. Consider:

London had snow yesterday. It fell to a depth of a metre.

What features must *London*, *snow* or *yesterday* be given to ensure that *it* is identified with *snow* and not with the other nouns?

When NL is used in a limited domain, semantic feature matching can be useful since a finite set of features may be adequate to account for the meaning of all relevant nouns. Outside such a domain, it seems unlikely that the number of features would be dramatically less than the number of words. For example, Fromkin & Rodman (1988, p.208) describe the word *mare* by features equivalent to [-HUMAN, -MALE, -YOUNG, +HORSENESS]. This enables us to distinguish:

<i>mare</i>	[-HUMAN, -MALE, +FEMALE, -YOUNG, +HORSENESS, -SMALL]
<i>stallion</i>	[-HUMAN, +MALE, -FEMALE, -YOUNG, +HORSENESS, -SMALL]
<i>foal</i>	[-HUMAN, ±MALE, ∓FEMALE, +YOUNG, +HORSENESS, +SMALL]
<i>pony</i>	[-HUMAN, ±MALE, ∓FEMALE, -YOUNG, +HORSENESS, +SMALL]
<i>horse</i>	[-HUMAN, ±MALE, ∓FEMALE, ±YOUNG, +HORSENESS, ±SMALL]

But what about *donkey*? Is this +HORSENESS with some other feature distinguishing a horse from a donkey? Or is it +DONKEYNESS, in which case how do we show the similarity between a horse and a donkey? Is *sincerity* to be classified as [-COUNT, +ABSTRACT, +SINCERENESS]? In which case why not just use [+SINCERENESS] with the rule:

+SINCERENESS ⇒ [-COUNT, +ABSTRACT]

Putting it another way, once nouns begin to require almost unique features, why not use the nouns themselves as the primitives?

In addition, the 'rules' may be broken deliberately, for rhetorical or poetical effect. Thus *The very stones cried out for revenge* breaks the normal restriction that *cry out* needs a subject with the feature +HUMAN, but is nevertheless perfectly understandable and thus meaningful in an appropriate context. Hence semantic features do not create ABSOLUTE rules.

4 θ-roles

Verbs need a somewhat different treatment. In the discussion above, to account for the oddity of:

**Ideas sleep.*
 **The stones cried.*

I claimed that *sleep* and *cry* require the 'subject' to be marked +ANIMATE. However, the term

‘subject’ is not adequate, since in a passive sentence,³ the grammatical subject is not the ‘semantic subject’. Thus both these sentences:

- **Sincerity puts the bottles on the table.*
- **Bottles are put on the table by sincerity.*

are semantically anomalous, even though *sincerity* is the subject in the first sentence (as shown by the agreement in number between *sincerity* and *puts*) whereas *bottles* is the subject in the second sentence (as shown by the corresponding agreement between *bottles* and *are*). Rather than analysing the sentence in terms of its syntactic components (e.g. subject, object), we should rather consider its semantic components. The verb *put* allows three semantically distinct constituents to be present in the sentence: the entity causing the action, the entity directly affected by the action and the location involved in the action.

Several detailed approaches to semantic analysis have been based on this insight. The core idea is basically that of ‘slot-and-filler’: a given verb is considered to generate a number of semantically defined ‘slots’ which then have to be filled by constituents of the sentence. ‘Case grammars’ and methods based on predicate logic are two examples of this core idea. Thus

The boy put the bottles on the table.

might be represented in predicate logic using the predicate *put* plus three arguments, e.g. *put(boy, bottles, on(table))*. I’ve chosen in this module to use a more directly linguistic account, based on Halliday. This talks of **thematic roles** or **theta roles** (θ -roles) in relation to the verb. The verb *put* allows three θ -roles: Agent, Patient and Location. In the sentence:

The boy put the bottles on the table.

the boy is the Agent (i.e. the doer of the action), *the bottles* is the Patient (i.e. the thing that suffers the action) and *on the table* is the Location. These phrases have exactly the same semantic roles in the passive sentence:

The bottles were put on the table by the boy.

although two of them have different syntactic roles. Using the passive form of the verb shows that the Agent θ -role can be omitted:

The bottles were put on the table.

although the other two θ -roles are compulsory. It is a general rule in English that the *by* PP in a passive sentence is optional.

An important distinction which is needed is between **Argument** and **Adjunct** θ -roles (also called **Participant** and **Circumstantial** θ -roles).⁴ Consider these two sentences:

- He was putting apples in the kitchen.*
- He was eating apples in the kitchen.*

In the kitchen is a Location in both sentences. However, it is an Argument only in the first sentence; in the second it is Adjunct. An Argument θ -role is one which is INTRINSIC to the action of the verb. Even when an Argument θ -role can be omitted from a sentence, it is still implied. For example, in passive sentences the Agent can be omitted, but we know that there must BE an Agent. By contrast, an Adjunct θ -role describes some EXTRA circumstances surrounding the main action. One test for the different types of θ -role is that an Adjunct θ -role can often be moved within a sentence without changing its meaning; in particular words

³ Action sentences in English can be classified as active or passive. Active sentences have a pattern like *The cat chased the mouse* or *The cat was chasing the mouse*. The first noun phrase in an active action sentence is the entity carrying out the action expressed by the verb. Passive sentences have a pattern like *The mouse was chased by the cat* or *The mouse was being chased by the cat*. The first noun phrase in a passive action sentence is the entity suffering the action expressed by the verb. Syntactically, the passive is formed by the auxiliary *be* plus the ‘en’ or ‘ed’ form of the verb, optionally followed by a prepositional phrase beginning with *by*.

⁴ Prior to 2007/08, I used Participant and Circumstantial, hence old examination papers ask about these terms. Now I prefer Argument and Adjunct.

like *while* or *when* can be inserted before it. Argument θ -roles tend to occur in fixed positions and cannot be moved without changing meaning. Thus:

He was putting apples in the kitchen.
 **He was in the kitchen putting apples.*
 **He was putting apples while in the kitchen.*
 **While in the kitchen, he was putting apples.*

He was eating apples in the kitchen.
He was in the kitchen eating apples.
He was eating apples while in the kitchen.
While in the kitchen, he was eating apples.

Given this distinction, the lexical entry for a verb must specify:

- The syntactically valid complements of the verb in an ACTIVE sentence. For *put* the only valid complements are NP+PP. The components of a verb complement (as defined so far) correspond to those Argument θ -roles which occur after the verb (in English).
- The semantically valid ARGUMENT θ -roles. For *put*, these are Agent, Patient and Location (ordered as they occur in an active sentence). Note that only Argument θ -roles need specifying (hence I will in future omit this qualification); optional Adjunct θ -roles can be added to any sentence as required.
- Whether the complements/ θ -roles are optional or not. Note that it is a general rule of English that the Agent θ -role is optional, so this does not need to be made explicit. (However since a grammatical subject is required in complete English sentences, the Agent θ -role can normally only be omitted in passive sentences, as noted in the general passivization rule given above.)

Given the discussion so far, a possible lexical entry for *put* is of the form:

put : complements = NP+PP, θ -roles = Agent:[+ANIMATE]+Patient+Location

The first θ -role corresponds to the subject NP in an active sentence, where it is automatically essential (but optional in the corresponding passive sentence). The remaining θ -roles correspond 1:1 with the components of the verb complement. It is perhaps clearer to extend the definition of verb complements to include the subject NP. A θ -role can then be associated with each component. Given this approach, the lexical entry for *put* might be:

put : complements = NP/Agent:[+ANIMATE] + NP/Patient + PP/Location

What this lexical entry means is that *put* should occur in sentences of the form:

NP/Agent:[+ANIMATE] *put* NP/Patient PP/Location

The syntax I'm using for a single complement is:

Nonterminal / Role : List_of_Features, where a list is constructed using [.]

The Role and List_of_Features are optional.

How can a PROGRAM use the lexical entry given above to analyse the sentence *The girl put the apples in the kitchen*?

- The sentence is parsed to yield $(the\ girl)_{NP} (put_V (the\ apples)_{NP} (in\ the\ kitchen)_{PP})_{VP}$, from which the verb is *put* with the full set of complements $(the\ girl)_{NP} + (the\ apples)_{NP} + in\ the\ kitchen)_{PP}$.
- Either during parsing or afterwards, program should attempt to match the complements one by one against the lexical entry for *put*, checking both the syntactic role and any features:
 - $(the\ girl)_{NP}$ matches NP/Agent:[+ANIMATE], provided that the lexicon for *girl* specifies +ANIMATE, so that Agent = *the girl*.
 - $(the\ apples)_{NP}$ matches NP/Patient, so that Patient = *the apples*.

- $(in\ the\ kitchen)_{PP}$ matches $PP/Location$, so that $Location = in\ the\ kitchen$.⁵

The process will reject sentences on both syntactical and semantic grounds. Thus *The girl put small in the kitchen* will be rejected because *small* is not an NP, whereas *Sincerity put apples in the kitchen* will be rejected because *sincerity* is [-ANIMATE]. If a sentence is accepted, then θ -roles will have been assigned. Although I am not going to consider how to code parsing in this way, it should be clear that it is only an extension of the methods discussed earlier.

Using this new definition of verb complements, we can write a rule to describe how a passive sentence in English is formed from an active one (a transformational grammar rule):

If and only if NP is the SECOND component of the lexically valid complements of an active verb, valid complements for the passive verb are formed from the rest of the verb complements plus an OPTIONAL PP with the preposition *by* (here written as PP(*by*)).

Semantic restrictions are also needed, so that the (almost) complete passivization rule is:

If and only if valid active verb complements have:

- NP/Agent as the first component
 - NP/<Role> as the second component
- the corresponding valid passive verb complements have:
- NP/<Role> as the first component
 - the remaining components of the active verb complements
 - optionally PP(*by*)/Agent as the final component.

Let's apply this to *put*, which was given the lexical entry:

put : complements = NP/Agent: [+ANIMATE] + NP/Patient + PP/Location

Matching this with the passivization rule above shows that *put* does have NP/Agent as the first component, and has NP/Patient as the second component, so that it fits the requirements for the rule, with <Role> = Patient. Hence the valid passive verb complements are:

- NP/Patient
- PP/Location
- optionally PP(*by*)/Agent: [+ANIMATE].

So given the sentence:

The apples were put in the kitchen by the girl.

the process of analysis should now be as follows.

- Parsing the sentence determines from the form of the verb that it is passive (basically because it has *were put* instead of *put*), with the structure $(the\ apples)_{NP} (were_{aux}\ put_V (in\ the\ kitchen)_{PP} (by\ the\ girl)_{PP})_{VP}$.
- The passive complements of *put* are NP/Patient + PP/Location + {PP(*by*)/Agent: [+ANIMATE]} (using the passivization rule on the active complements in the lexicon).

Hence matching complements one by one gives Patient = *the apples*, Location = *in the kitchen* and Agent = *the girl* (assuming *girl* has the lexical entry +ANIMATE).

Using the same approach, a program can reject:

Girls were put in the kitchen by the apples.

since the supposed Agent, *apples*, is not animate so $(by\ the\ apples)_{PP}$ doesn't match the expected complement PP(*by*)/Agent: [+ANIMATE].

My conclusion is that it is possible to computerize at least some processing of the meaning of simple sentences such as these. Semantic features (e.g. [+ANIMATE]) can be used to handle the meaning of nouns and adjectives, with θ -roles used to process the meaning of verbs plus

⁵ Typically more checks would be made here on the match, since only certain English prepositions can specify a Location, e.g. *in*, *on*, *outside*.

the meaning associated with some of the word ordering within sentences. To the extent that semantics was defined as ‘compositional meaning’, i.e. meaning based on the meaning of words plus meaning associated with the ordering of words within sentences, then limited semantic processing has been achieved.

Further stages might involve representing the ‘extracted meaning’ using some form of logical calculus. For example, having analysed *The apples were put in the kitchen by the girl* in this way, the meaning could be represented in the predicate calculus, e.g. as

$$\exists A \exists P \exists L (\text{girl}(A) \ \& \ \text{apple}(P) \ \& \ \text{kitchen}(L) \ \& \\ \text{put}(\text{agent}(A), \text{patient}(P), \text{location}(L)))$$

This allows further reasoning, e.g. by combining with a general rule such as

$$\forall A \forall P \forall L (\text{put}(\text{agent}(A), \text{patient}(P), \text{location}(L)) \rightarrow \text{in}(P, L))$$

we can deduce

$$\exists P \exists L (\text{apple}(P) \ \& \ \text{kitchen}(L) \ \& \ \text{in}(P, L))$$

i.e. *the apples are in the kitchen*.⁶

5 Case Study: *break*

A case study of the verb *break* can demonstrate some of the complexity and subtlety of the approach outlined above. In its ‘physical’ sense, *break* could have a lexical entry of the form:

$$\textit{break} : \text{complements} = \text{NP/Agent: [+ANIMATE]} + \text{NP/Patient: [-ANIMATE]} \mid \\ \text{NP/Patient: [-ANIMATE]}$$

The lexical entry states that when the verb complements are NP+NP, the first NP must be an animate Agent and the second an inanimate Patient. The passivization rule shows that sentences with these complements can be made passive, the passive complements being NP/Patient: [-ANIMATE] and optionally PP(*by*)/Agent: [+ANIMATE]. Thus from the first set of complements we have:

The man broke the cup.

The cup was broken by the man.

The cup was broken. (Agent is optional in passive sentence.)

**The cup broke the man.* (Agent is not animate, Patient is not inanimate.)

**The man was broken by the cup.* (Patient is not inanimate, Agent is not animate.)

**The man was broken.* (Patient is not inanimate.)

When the verb complement is just NP, it must be the Patient and passive sentences cannot be formed because the passivization rule does not apply. Thus:

The cup broke.

**The man broke.* (Patient is not inanimate.)

Another possible θ -role is Instrument – the object used to achieve the action. For example, in the sentence:

The man broke my cup with his hammer.

the man is the Agent, *my cup* is the Patient and *his hammer* is the Instrument. The preposition *with* usually introduces the Instrument in English. Instrument is optional, so that the verb complements for *break* can be re-written as:

$$\textit{break} : \text{complements} = \text{NP/Agent: [+ANIMATE]} + \text{NP/Patient: [-ANIMATE]} + \\ \{ \text{PP/Instrument: [-ANIMATE]} \} \mid \\ \text{NP/Patient: [-ANIMATE]}$$

The passivization rule gives the corresponding verb complements in a passive sentence as NP/Patient: [-ANIMATE] + {PP/Instrument: [-ANIMATE]} + {PP(*by*)/Agent: [+ANIMATE]}. Thus it correctly predicts the validity/invalidity of the following sentences:

The cup was broken with the hammer by the man.

⁶ More accurately in the representation here *Some apples are in a kitchen*.

I hope that exploring this one verb in considerable depth has shown (a) the complexity of analysis which may be needed (b) the ability of θ -roles to explain subtle semantic differences. The Exercises look at some other verbs and θ -roles.

6 Classification of Verbs

Verbs can be classified using the θ -roles they involve. Dictionaries often use the term ‘transitive’ to describe a verb which has at least two NP θ -roles (usually Agent and Patient), and ‘intransitive’ to describe a verb which has only one NP θ -role (in English either Agent or Patient). This two-fold classification, however traditional, is clearly inadequate, and the full range of θ -roles must be considered.

Two further θ -roles are Source (where the action originated) and Goal (where the action is directed). In the sentence:

The woman carried the box from her house to the car.

the woman is the Agent, *the box* the Patient, *her house* the Source, and *the car* the Goal. Particular prepositions often identify the Source (e.g. *from* or *out of*⁷) and the Goal (e.g. *to* or *into*).⁸ The verbs *carry*, *drag*, *haul* or *push* all seem to have the same θ -roles (Agent, Patient, Source, Goal), with the same rules about optionality (Patient is essential, Goal and Source are optional). Thus we have the examples:

I carried the box from the house to the car.

I dragged the box to the car.

*I hauled the box out of the house.*⁹

I pushed the box.

**I carried.* (Patient may not be omitted.)

It is, presumably, no accident that these verbs have the same θ -roles, since they have ‘similar meanings’. Schank and co-workers, from 1973 onwards, developed a theory called Conceptual Dependency (CD) which focussed on verb-based semantic primitives. Each primitive has a particular set of θ -roles (CD terminology is somewhat different, but the principle is the same); more importantly, each represents a ‘basic action’. *Carry*, *drag*, *haul* and *push* are all examples of the primitive PROPEL, which involves moving something by the application of physical force. Proponents of CD hoped to find a minimum set of such primitives into which any sentence could be decomposed. The primitives could then be used in further reasoning.

In essence, this approach is feature-based: *carry*, *drag*, *haul* and *push* are being given the feature +PROPEL, just as *man*, *woman*, *boy* and *girl* were earlier given the feature +HUMAN. The feature +PROPEL carries with it a particular set of θ -roles. CD thus suffers from the same problem as other feature-based approaches. If the classification used is too general, then differences in meaning will be missed. For example, in CD theory, *smoke* was classified with *eat* and *drink* as an example of the primitive INGEST, but these actions differ significantly. (For example, if you eat or drink something, every part of it ends up inside you, which is not the case when smoking a cigarette.) On the other hand, if the classification is too specific, little is gained from using it in preference to the original verbs. CD theory is not generally used today, and it seems that no method of classifying verbs via intrinsic semantic features has stood the test of time.

However, a weaker classification based solely on θ -roles is useful, if only to reduce repetition in the lexicon. Part of such a classification might be:

⁷ Although *out of* is written as two separate words, it is usually treated as a single preposition (cf. *into* written as one word but derived from *in to*).

⁸ In other languages, inflections are used to mark θ -roles, with or without prepositions. A good example is the non-Indo-European language Finnish.

⁹ At least in my idiolect, *from* is not normally acceptable without *to*. Thus I can say *I carried the box out of the house* but not *I carried the box from the house*. *From .. to ..* behaves in effect as a unit. Whether this is a syntactic or a semantic constraint isn’t clear to me.

Action-Intransitive

complements = NP/Agent: [+ANIMATE]

Action-Transitive

complements = NP/Agent: [+ANIMATE] + NP/Patient + {PP/Instrument: [-ANIMATE]}

Action-Transitive-And-Intransitive

complements = NP/Patient: [+ANIMATE]

Action-Bitransitive-To

complements = NP/Agent: [+ANIMATE] + NP/Patient + {PP(*to*)/Recipient}

Action-Bitransitive

complements = NP/Agent: [+ANIMATE] + NP/Recipient + NP/Patient

Categories are hierarchical, inheriting complements from those above. Thus *clean* is Action-Transitive:

Joe cleaned the floor.

**The floor cleaned.*

whereas *dry* is Action-Transitive-And-Intransitive, a subclass of Action-Transitive:

Joe dried the floor.

The floor dried.

Similarly *donate* is Action-Bitransitive-To, *give* is in the subclass Action-Bitransitive:

Ravi donated the money to Oxfam.

**Ravi donated Oxfam the money.*

Ravi gave the money to Oxfam.

Ravi gave Oxfam the money.

Verbs can belong to more than one category. For example, *cry* (in the sense of ‘weep’) is both Action-Intransitive as in:

John cried.

and Action-Transitive as in:

John cried tears of joy.

Restrictions on allowed verb complements, like those based on semantic features, are not absolute. In ‘appropriate’ contexts otherwise mandatory θ -roles can be omitted. Thus:

**I give.*

I give and he takes. (Idiom?)

**I put the forks.*

He put the knives on the table and I put the forks. (Ellipsis – see below)

7 θ -roles in Machine Translation

Identifying θ -roles seems to be important in translation. English tends to be very flexible (not to say vague) in relating syntactic and thematic roles. Thus in an active sentence involving *break*, the grammatical subject can be the Agent, Force or Patient:

The boy broke the window.

The wind broke the window.

The window broke.

The French verb *casser* (meaning ‘to break’) allows only the Agent or Force roles as the grammatical subject, giving:

Le garçon a cassé la fenêtre. (*a cassé* ♠ has broken; French has no ‘Simple Past’)

Le vent a cassé la fenêtre.

**La fenêtre a cassé.*

The semantic equivalent of *The window broke* is possibly:

La fenêtre s’est cassée. (*s’* = *se*)

which could be translated as ‘the window broke itself’, semantically equivalent to the passive ‘the window was broken by itself’. The window is thus both Force AND Patient, which is perhaps clearer than the equivalent English sentence. The passive without an Agent is less common in French, so that an alternative translation of:

The window was broken.

is:

On a cassé la fenêtre.

literally ‘Someone broke the window’.¹⁰

Notice the correspondence of the θ -roles in these examples. *The window was broken* is derived from a passive transformation of the verb complements NP/Agent + NP/Patient, with Agent optionally omitted. Hence we know that there is an Agent, although not who it is. The French *On a cassé la fenêtre* conveys the same information: *on* is a non-specific Agent – there is an Agent, but we do not know who it is. On the other hand, in *The window broke*, there is no Agent, so that if I ask *Who broke it?* an appropriate answer would be *I don’t know, it just broke*. Similarly in the French *La fenêtre s’est cassée* there is no external Agent: the window is both Force and Patient; it just broke itself, the sentence implies.

This suggests a different model of sentence translation. By a combination of syntactic and semantic processing of the first language, we determine the ‘action’ (verb) plus the associated θ -roles. These are then translated into the second language, possibly using entirely different syntax, but retaining the same action and θ -roles.

So far I have only considered one semantic class of sentence, so-called ‘Action Sentences’ where the verb involves something being done to or happening to something. Now consider the English sentence:

I liked that book.

I does not seem to be in the Agent role, nor does *that book* seem to be in the Patient role, in spite of their positions in the sentence: I do not do anything to something by liking it. When a sentence has clear Agent and Patient θ -roles, there is a paraphrase of the form *What <Agent> did to <Patient> was to <verb> it*. Thus we have:

I broke the cup.

What I did to the cup was to break it.

BUT:

**What I did to that book was to like it.*

If there is a paraphrase of this form it is:

What that book did to me was to cause me to like it.

Possible names for the new θ -roles involved are Processor and Phenomenon. *I* is the Processor, *the book* the Phenomenon. The Processor does not do anything to the Phenomenon, which in turn is not altered or affected in any way by being processed. Verbs involving ‘mental processes’ such as *like, see, hear*, etc. seem to involve these new θ -roles.

English, as we noted above, tends to have limited syntactic marking of different θ -roles compared to some other languages. Hence it is not surprising to find that Spanish and Greek translations of *like*, for example, require *I* to be in a clearly non-Agent θ -role:

Me gusta este libro (Spanish) = me pleases this book = This book pleases me

Mou aresi avto to vivlio (Greek) = of-me pleases this the book = This book pleases me

This example again reinforces the idea that something like θ -roles are needed as the intermediate in translation. The problem is to decide whether there is a finite set of distinct θ -roles (preferably marked with appropriate semantic features). In a restricted domain, this answer is likely to be ‘yes’. For example, in translating maintenance manuals or weather forecasts, it is

¹⁰ But *on* is a much vaguer pronoun than *someone* since it can stand for ‘I’ or ‘we’, as in *on y va?* meaning ‘shall we go?’ Note also that in French the verb *aller* requires a Goal (*y* is a kind of ‘pro-Goal’, a non-specific place, something like the English *there* but even less specific). Hence **on va?* is unacceptable.

reasonable to assume that verbs will be used in specific ways with clearly defined θ -roles. Thus in both contexts we can assume that *break* will have an inanimate Patient as in:

Overtightening may break the glass.
The spell of hot weather was broken by today's heavy rainfall.

rather than an animate Patient as in:

Three years of solitary confinement finally broke him.

How far this can be extended to NL in general is less clear. If we have to invent a large number of θ -roles, some of which apply to only a few verbs, we are back to the equivalent of +HORSENESS or +SINCERENESS and we might as well have meaning/translation rules for each individual verb.

8 Postscript

I have only been able to skim the surface of semantic processing and have not talked about implementation. θ -roles have not been used directly in many NLP systems, although 'case grammars', which incorporate many of the same ideas, have, and both are examples of the 'slot-and-filler' approach whereby the lexicon specifies an 'schema' associated with a lexeme and processing involves discovering components to fill slots in the schema. MANY more issues need considering before a practical system can be constructed. These include the following.

- **Anaphora.** Consider:

*There's a man at the door. **He** wants to speak to you.*
*I wanted to finish today, but I couldn't **do it**.*
*I put it on the table. Is it still **there**?*

The emboldened word(s) in each case refer back to some component which appeared earlier. *He* refers to *a man*; *do it* refers to *finish today*; *there* refers to *on the table*. The process of 'reference back' is called 'anaphora'. All languages allow various kinds of anaphora.

In semantic processing we need to be able to identify the referent. How can this be done? In some cases semantic features may help. In the first example above, *man* is [+ANIMATE, +HUMAN, +MALE] as is *he*, so it is reasonable to equate them. But what features are needed to identify *it* correctly in the example I gave earlier:

*London had snow yesterday. **It** fell to a depth of a metre.*

- **Ellipsis.** Closely connected with anaphora is ellipsis: the omission of components of sentences which can be deduced from earlier occurrences. Thus:

Mary likes Tom but hates William. (i.e. ... but Mary hates William.)
I've finished cleaning. Have you? (i.e. Have you finished cleaning?)
I wanted to finish it today, but couldn't. (i.e. ... but I couldn't finish it today.)

Notice that in some cases, inserting the appropriate 'pro-form' converts the ellipsis to anaphora, e.g.:

*Mary likes Tom but **she** hates William.*
*I wanted to finish it today, but couldn't **do it**.*

- **Presupposition.** The meaning of some sentences crucially depends on implicit presuppositions. The most notorious English example is probably:

Have you stopped beating your wife?

which presupposes:

You beat your wife.

so that neither a *Yes* or a *No* answer is possible if in fact you never beat your wife. (It also presupposes that you have a wife, which in turn presupposes that you are male.) Anaphora normally presupposes that the hearer knows the reference. Thus if I say:

Give him the book.

I am assuming that the person to whom I am speaking knows who *him* refers to.

More subtle examples of presupposition show up in sentences like:

I gave the money to John.

I gave John the money.

Using only θ -roles, these sentences appear to have the same meaning: the Action is *give* (in the past), the Agent is *I*, the Patient is *the money* and the Recipient is *John*. However, spoken with normal stress and intonation, the first sentence assumes that the hearer knows that I gave money to someone, but not to whom, whereas the second sentence assumes that the hearer knows that I gave something to John, but not what. It is for this reason that we have:

I didn't want the money, so I gave it to John.

**I didn't want the money, so I gave John it.*

For John's birthday, I gave him a tie.

**For John's birthday, I gave a tie to him.*

The Agent/Force distinction, noted above, also generates important presuppositions. Consider:

We were playing football and the ball broke the window.

According to the analysis given above, this puts *the ball* in the Force θ -role, which in turn presupposes that no Agent exists, so that no-one was responsible for the action of breaking the window.

- **Discourse rules.** A connected text or a conversation between people is acceptable only if the appropriate discourse rules are followed. Some are non-semantic, such as turn-taking, others are semantic, such as relevance maintenance. One rule is that 'known' information should be de-stressed and 'new' information stressed, presumably to make it easier for the hearer to pick out the important points. In English, new information generally goes as far to the end of the sentence as is syntactically possible. Anaphora, ellipsis and presupposition are often used to remove 'interfering' material. Consider this conversation:

A: *I didn't want to look after the money.*

B: *So you gave it to John?*

A: *Yes, I trusted him.*

B: *But you didn't get it back?*

A: *No.*

Filling out the anaphora, ellipsis and presupposition might give:

A: *I didn't want to look after the money which you already know about.*¹¹

B: *You gave the money to John because you didn't want to look after the money?*

A: *I gave the money to John because I trusted John.*

B: *You didn't get the money back after you had given the money to John?*

A: *I didn't get the money back.*

People prefer the first version of the conversation. However, given our present state of knowledge, it seems that a computer program could more easily be written to process the second version.

- **Pragmatics.** We ought also to consider 'pragmatics', i.e. take into account the INTENTIONS of the NL user. The boundary between semantics and pragmatics is fuzzy, so an example may clarify my usage.

Can you hold it still?

Semantically this is a question, eliciting information. A semantic paraphrase would be:

Is it the case that you are able hold it still?

¹¹ The word *the* in the original signals that the money is not a new topic. Consider:
A: *There's a man at the door.* (*a* = new information + numerically one)
B: *Tell the man to come in.* (*the* = not new information, hence a reference back)

If a doctor said the original sentence to a patient suffering from a motor disease, this would probably be the correct interpretation. However, in ‘normal’ English use, a more likely semantic paraphrase is:

Hold it still, please.

NL is frequently not intended to inform but to persuade or even mis-inform. ‘Smuggling in’ erroneous presuppositions and implications is one way of doing this. Compare:

We were playing basketball and the ball broke the window.

Some friends and I were playing basketball and the window was broken with the ball.

The ‘facts’ are the same: a number of people, including the speaker, were playing basketball. The window is now broken. The ball was involved. The differences are that the use of *we* in the first sentence de-emphasizes the involvement of the speaker compared to *some friends and I*, and the use of *the ball broke the window* implies that no Agent was involved, whereas *the window was broken with the ball* implies that there was an Agent involved (which could well have been the speaker). ‘Filtering out’ the speaker’s intention to mis-inform is an vital human skill!

Exercises

1. In the text, three rules were suggested for the relationships among the features COMMON, COUNT, ANIMATE, HUMAN, MALE, FEMALE, ABSTRACT. Suggest THREE more relationships.
2. Consider the grammatical subject NP in these sentences:

Jane is a person.

**The Jane is a person.*

Everest is a mountain.

**The Everest is a mountain.*

**Himalayas are mountains.*

The Himalayas are mountains.

**Thames is a river.*

The Thames is a river.

London is a city.

**The London is a city.*

England is a country.

**The England is a country.*

**United States is a country.*

The United States is a country.

Can you formulate rule(s) based on semantic features to account for the distribution of the determiner *the*? That is, can you suggest values of `featureList1` and `featureList2` to allow these noun phrases to be generated correctly by a grammar of the form:

`NP` → `det SNP(featureList1) | SNP(featureList2)`

`SNP(FL)` → `adj SNP(FL) | noun(FL)`

`noun(FL)` → {any word in the lexicon as a noun and with a list of features including those in FL}

Himalayas : noun, [features including those in `featureList1`]

Jane : noun, [features including those in `featureList2`]

3. *All*, *most*, *much* and *less* are some of the English ‘quantifiers’. Consider the grammatical subject NP in these sentences:

All boys are horrid.

Most boys are horrid.

**All boy is horrid.*

**Most boy is horrid.*

All sincerity is faked.

Most sincerity is faked.

**Much boys are horrid.*

**Less boys are horrid.*

**Much boy is horrid.*

**Less boy is horrid.*

Much sincerity is faked.

Less sincerity is faked (than ...)

Can you formulate rule(s) based on semantic features to account for the distribution of the quantifiers? Think of THREE more English quantifiers. Do they follow the same rule(s)? Try to complete the grammar starting:

`NP` → `quant(..) SNP(..)`

4. Chomsky’s notorious sentence *Colourless green ideas sleep furiously* is highly semantically anomalous. Suggest semantic features for all five words which demonstrate all the anomalies.

5. a) Suggest a suitable lexical entry, of the form given earlier, for the English verb *eat*. Construct sentences, both active and passive where appropriate, to illustrate correct and incorrect sentences using this verb, showing IN DETAIL, as I did for *put* above, how their analysis follows from your lexical entry.
- b) Consider sentences like *The boy opened the door with a key*, *The key opened the door* and *The door opened*. Repeat (a) for the verb *open*. Does *open* have the same verb complements as *break*? (The two verbs seem to share some semantic similarity.)
- c) Repeat (a) for the verb *give*. You will need a new θ -role, Recipient.
- d) The verb *throw* seems to have similarities to both *push* and *give*. Thus:

He threw it from the house to the garden. *He pushed it from the house to the garden.*

The boy threw the ball into the house. *The boy pushed the box into the house.*

The boy threw the ball to Jo. *The boy gave the ball to Jo.*

The boy threw Jo the ball. *The boy gave Jo the ball.*

Construct other sentences OF THIS KIND, both active and passive where appropriate, to illustrate correct and incorrect sentences using *throw*. Do you agree that its lexical entry should be a combination of those for *push* and *give*?

- e) Verbs like *drive* and *throw* (less so *push*) can be used with a different preposition, namely *at*:

I drove my car at the wall.

The boy threw stones at the cat.

?*He pushed the gun at the policeman.* (*thrust* is better here.)

Is this one of the θ -roles introduced so far, or is it a new one? Present evidence for whichever decision you make.

- f) If you know another language well, or can find some willing speaker of another language, analyse how some of the sentences you considered in (a) to (e) should be translated. Is there evidence that θ -roles are more important than syntax?
-