

Yorick Wilks  
Dept. of Language & Linguistics,  
University of Essex,  
Colchester,  
ENGLAND.

### Abstract

The paper discusses the incorporation of richer semantic structures into the Preference Semantics system: they are called pseudo-texts and capture something of the information expressed in one type of frame proposed by Minsky (q.v.). However, they are in a format, and subject to rules of inference, consistent with earlier accounts of this system of language analysis and understanding. Their use is discussed in connection with the phenomenon of extended use: sentences where the semantic preferences are broken. It is argued that such situations are the norm and not the exception in normal language use, and that a language understanding system must give some general treatment of them.

Descriptive terms: natural language understanding, semantics, frames, pseudo-text, preferences, templates, formulas, thesaurus.

### Introduction

This paper sketches how one might deal with extensions of word sense in a natural language understanding system (NLUS): that is to say, normal utterances that break preassigned semantic selection, or preference, restrictions. The proposals here extend the knowledge representation of the preference semantics NLUS (Wilks 1968, 1975) with pseudo-texts (PT) which are one type of frame structures in the sense of (Minsky 1975), but which are also consistent with the general assumptions of this particular NLUS. At the end I shall describe an implementation environment under construction, which may make possible some test of the relative contributions of PTs and very general pragmatic principles like "preference" (see below) to language extension.

To understand "preference" consider the following sentence, chosen, I promise you, at random from the front page of a daily newspaper: (The Times 5.2.76):

(1) Mr. Wilson said that the line taken by the Shadow Cabinet, that an Assembly should be given no executive powers would lead to the break-up of the United Kingdom.

This sentence presents no problems to the average reader of that newspaper, who is presumed to know what a cabinet is, and what the U.K. is. However, at each of the four underlined points, the noun would violate the normal semantic selection restrictions for the associated verb: lines, for example, would violate the normal "physical

object" restriction on "take" and so on.

I shall refer to such restrictions as preference restrictions, because of the way the present NLUS is already able to accept natural language that violates preferences, as (1) does (see recap in next section for more detail). Such usage as (s) will be referred to as extended, or preference violating, and these will serve instead of the more literary and philosophical term "metaphorical".

It is an important assumption of this paper that such usage is the norm in ordinary everyday language use, and cannot be relegated to the realm of the exceptional, or the odd, and so dealt with by considerations of "performance". On the contrary it is, I would argue, central to our language capabilities, and any theory of language must have something concrete to say about it. Even if the newspaper usages above are "extended", I would suggest that anyone who could not grasp these extension could not be said to understand English properly (given adequate knowledge from which to extend, and we shall come to that.) It will be obvious already that the commitment to a norm implies a corresponding commitment to general everyday language as a proper topic for AI. This assumption needs defence, but there is no space for it here. However, one might bear in mind that, although non-general micro-worlds have been put forward as the E.Coli et Al, they may in fact turn out to be our phlogiston!

### Semantic Sense-Projection

The process described in this paper is called projection: sense descriptions for words will be rewritten, in preference-violating texts, with the aid of the specific knowledge in PTs; which is to say that part of the PT is projected into the sense description for a word. So, for example in (1) some detailed political knowledge in a PT (see below) for "United Kingdom" could show that a breaking of such an entity could be caused, and we would then replace the sense description of "lead to the break-up" and providing a more appropriate sense description of "lead to" for analysis of the rest of this text.


The essence of this process is finding in the appropriate PT, what it is that can normally be done with an entity of that type. As we shall see, matching this is not always a simple matter. If, with the aid of knowledge representations and sense descriptions, projection, even to a small degree, and generalized manner, we shall, at the same time, be able to explain why the same acceptable extenders use is not present in, say, "He broke up evil",

Thus the present task of explication will, to that extent, make concessions to the main linguistic goal of the last two decades, that of setting boundaries to acceptability or meaningfulness, at least in so far as (1) would be interpretable on the basis of a knowledge base, and rules of extension, while the last sentence would not.

However,, some brief recap of the existing state of the NLUS is necessary for setting out these extensions.

#### Brief recap of preference semantics

In previous papers I have described an NLUS in which rules operate on semantic word-sense descriptions to build up text descriptions. The tules that insert sense descriptions into text descriptions are what I called "preferential": they seek preferred entities, but will accept the less preferred if necessary. A sense description for the action "drink" would be the semantic formula:



This is a formal structure of semantic primitives expressing the meaning of the action (see King and Wilks 1977): that drinking is a CAUSING to MOVE, preferably done by ANimate SUBject (=agent) and to a liquid (FLOW STUFF). TO a particular ANimate aperture (THRU PART), and INTO the SELF (=the animate agent). The last primitive cause, is the head of the formula and its main primitive. For short we will write (2) as |drink|.

The text structures in the system are templates together with semantic ties between them, where a template is a network of formulas, containing at least an agent, action and object formula (or appropriate dummies in place of them). Thus the template for "The adder drinks water" will be written for short [the+adder drinks water] in which (2) is at the (action) node.

The process of setting up the templates allows the formulas to compete to fill noes in templates. Thus tge formula for the (snake-)adder goes to the agent node in the template above in preference to the (machine-)adder because (2) specifies, by (\*ANI SUBJ) that it prefers to be accompanied in a template by an animate agent formula. However, in the sentence:

(3) My car drinks gasoline.

the available formula for the first template node namely [car], is not for an animate entity, yet it is accepted because there is no competitor for the position. THE PURPOSE OF THIS PAPER IS TO INVEST-GATE HOW THE SYSTEM MIGHT NOT MERELY ACCEPT SUCH A PREFERENCE-VIOLATING STRUCTURE FOR(3) BUT MIGHT ALSO INTERPRET IT.

An important later process is called extrac-tion: template-like structures are inferred and added to the text representation even though they match nothing in the surface text. They are "deeper" inferences from the case structure of formulas in

some actual template--where the case primitives are those underlined in (2). Thus, to the template for (3), we would add an extraction (in double square parentheses in abbreviated form):

(4) [[gasoline in car]]

which is an inference extracted from the contain-ment subformula of (2) (SELF IN). Analogous extractions could be made for each case primitive in each formula in the template for (3).

All these are, of course, complex and content-ous issues, that can only be summarised here so that we can get on to something else. They have however been programmed and described in detail in (Wilks 1975, 1976).


Since then a structural change (Wilks (1976a) as allowed a wider, and more specific, form of expression in formulas by allowing thesaurus items, as well as primitives, to function in them. No problems are introduced by doing this, provided that the thesaurus items are also themselves words in the dictionary, and so have theirformulas defined elsewhere in their turn. One advantage of this extension is to impose a thesaurus struc-ture on the whole vocabulary, and so render its semantic expression more consistent.

A thesaurus, Like Roget, is simply an organi-sation of a vocabulary into semi-synonymous rows, which are themselves classified hierarchically under heads, and even more generally, sections. Thus under some very general section name MOVE (=motion) we would find heads, two of which might be #engine and #vehicle. The former might be the name of a row of actual types of engine

(S) # 525 engine: turbine, internal com-bustion, steam.....

where the number simply indicates the sequence position of # engine in the thesaurus. It is no accident that the most general section names like MOVE can be identified with the semantic primitives of the present system.


The organisation is imposed by requiring inclusion relations, between the formulas for word senses, corresponding to the thesaurus relations of the words. Thus, all the words in the row (5) would have a common subpart to their formulas, and that common subpart would be the dictionary formula for "engine", probable expressing in primitives no more than "a thing used by humans to perform some task, and self-moving in some way". If now thesaurus items can be inserted in formulas we may expect a formula for "car" at least as specific as:



## Language Boundaries and Projection

Let us return to examples like (3) for which the system constructs a template even though it contains a violated preference, and ask what should an intelligent system infer in such a situation?\* I would suggest that cars can be said to drink in virtue of something a system might already know about them, namely that they have a fluid (gas/petrol) injected into them, and they use that in order to run. That is to say, the program should have access to a sufficiently rich knowledge structure to express the fact that cars stand in a relation to a particular fluid, a relation that is, of the "same semantic structure" as the relation in which a drinker normally stands to the thing drunk. All that may sound obvious, but how else are we to account for the naturalness of (3), but the relative unnaturalness (and uninterpretability) of "My car chews gasoline", and, the more distant, "My car carves the Sunday roast". One upshot of these proposals is to distinguish plausible (with respect to a knowledge base) preference violation from the implausible.\*\*

The procedural upshot of the above would be to replace at least one formula in the template for (3) with another, either constructed by rule\*\*\* or drawn from the knowledge structure itself, to be called a pseudo-text (PT). Let us now postulate that "car" points not only to (6), i.e. [car] but that [car] in turn points to :



Part of the pseudo-text for "car"

\*\* An important aspect of the interpretation of (3) is idiomatic, namely that the car uses a lot of gas/petrol. This aspect of the meaning is beyond this, or I suspect any, general inference procedure.

\*\*\* In a fuller version of this paper (Wilks, in press) I describe the relation of this work to attempts, such as (Givon 1967), to give general rules for projection: rules operating on the dictionary and independent of contexts of use.

This structure is called a pseudo-text because it is of just the same format as the text representations produced by the present NLUS. It can be extended to taste to express as much specific information about cars as is thought appropriate. Given the parser for the present NLUS, it could even be input as real text about cars. The representation consists of the templates (explained loosely at the right), together with the self-explanatory case and cause ties between them. In the templates, # denotes a dummy and \* denotes the formula [car] that points to the object (7). The # prefixed items are thesaurus items, though the "IC engine" is simply a specific dictionary word pointing to its own formula - specificity is thus a matter of taste. So, for example, the thesaurus head # liquid could be replaced by more explicit "gasoline". Items in round parentheses remain in primitive form.

It will be clear that the same information can be expressed in a number of different ways, and at different levels of generality; though the spirit of (Minsky 1975) suggests that they should be as specific as possible. The intention here is that THE PROCESSES THAT OPERATE ON SUCH ENTITIES AS (7) SHALL BE IDENTICAL WITH THOSE THAT OPERATE ON SUCH ENTITIES AS (7) SHALL BE IDENTICAL WITH THOSE THAT MANIPULATE REPRESENTATIONS DERIVED FROM INPUT TEXTS. The approach is thus the reverse of the conventional one: we seek to assimilate knowledge structures to text structures, rather than the reverse, on the grounds that the representation of language is the difficult task, and that the representation of knowledge as such makes no sense apart from that.

We should note too, that just as the thesaurus structure imposes a containment relation on the formulas of co-row-member words, so it also imposes a hierarchical relationship on PTs: that for # vehicle, for example, will be a less specific version of (7). Further up the thesaurus would be PTs for high-level sections: that for the primitive MAN would be highly complex. But note there is no "inheritance of property" problem in this system: the formula for "amputee" would have head MAN and would specify the loss of limbs. Any inherited pseudo-text from MAN-asserting "two legs" - would be modified [amputee].

The system now uses (7) to make a projection, so as to derive an interpretation for (4), by seeking, in (7), templates matching the source

template [my+car drinks gasoline]: namely the first and fourth lines of (7). The first match is in virtue of the similarity of [drink] and [# inject]--based on the expression in primitives, as in (2), of causing liquid to be in an entity of the same type as the agent. This would allow us to confirm, by projection, the "humanness of the drinker", that has already been noted by earlier extraction\* routines that extracted out from [drink] independently of the PT (7). However, no projection is made at this stage onto [car] (though it might be later in the face of a sentence like "His thirst is never slaked following (4), that confirms the humanness projection) because, in the case of violations of the preferences of actions, as of "drink" in (4), the system ALWAYS PREFERS TO MAKE A PROJECTION ON TO THE ACTION ITSELF IF IT CAN.

A stronger match is detected between the [my+car drinks gasoline] and the fourth line of (7) in virtue of the containment of [#engine] in [car], and of [gasoline] in [#liquid], which is evident in the formulas themselves. This results in the projection of the action node of the fourth line of (7), namely [use], on to [drink] in the template fir (3). This projection is taken to be strongly confirmed by the earlier match with the first line of (7), and is considered to carry over more sense than any alternative projection. The confirmation (of the match of the template to the fourth line of (7) by that of the first line) is necessary here, because [my+car leaks gasoline] would also match the fourth line, but no such projection would be appropriate. Conversely, no projection could be made for "My car drinks mud" from the fourth line, even with the confirmation of the first. The general rule for action projections then is: SEEK A PSEUDO-TEXT, FOR AGENT OR OBJECT, WITH A TEMPLATE MATCHING ON AGENT AND OBJECT NODES. PROJECT THIS GENERALLY IF THERE IS ALSO A PSEUDO-TEXT TEMPLATE MATCH TO THE ACTION ITSELF, FOR ANOTHER TEMPLATE IN THE SAME PSEUDO-TEXT.

We may note in brief because of pressure of space, above suggestion. First consider the more complex example presented by a recent headline:

(8) United Kingdom tries to escape Common Market. Clearly, some projection would be appropriate here, of humanness on to the country (which would require a MAN-head for [U.K.]; [Common Market]). The knowledge required might be drawn from [escape] alone, by simple extraction and without recourse to the pseudo-texts for either of the entities of the U.K. joining, but not of leaving the Common Market. In such circumstances more historical facts are not enough, even when highly structured. We might conceivably be able to project some notion [disassociate] onto [escape], from the U.K. PT

\* extractions, it will be seen, differ from projections in that they produce new template-like entities, rather than, as here, replacing formulas inside existing templates.

given some more sophisticated matching criterion that placed relevance above negation in such cases (i.e. would match [escape] with [associate] or [join].

Secondly, we might consider the problems presented by an example like:

(9) I see what you mean. Here the last clause breaks the preference expressed in [see] for a physical object. A system procedure will present the actual object of (9) to the top-level template simply as the primitive SIGN (the primitive for symbols and intensional representations of them) which has been obtained, by extraction, from the preferred object in [mean]. Thus the system is effectively dealing with the template sequence. [I see (SIGN)] [you man (SIGN)]. But what could we expect as a pseudo-text for something as general as SIGN, so as to use the above procedures to project on to [see]. If we take advantage of the hierarchical nature of the thesaurus, we might expect pseudo-texts at the very top level, associated with the section names - pure primitives like SIGN - just as specific pseudo-texts are associated with the lowest level items in the thesaurus - row members like "car". The pseudo-text for a primitive like SIGN would be wholly "care structural": it would consist of no more than primitive concatenations, in template form, like MAN THINK SIGN\*, the most general thing that can be said about what is normally done to signs. However, even something as general as this might suffice to project THINK correctly onto [see]. The interesting generality would come from using exactly the same projection procedures on the most general pseudo-texts like this, as on the most specific, like (7).

Thirdly, and this is treated at length in Wilks (1977a), we can consider a quite different type of projection for phrases like (10) a toy lion.

This comes from a much discussed class of examples ("plastic flower", "stone horse" etc.), where an obvious projection mechanism is to replace the head of the formula for the noun (BEAST) in [lion] in (10) by the preferred object of predication in the qualifier -- here \*PHYSOB in [toy]. This would be a very limited and general class of projections, not requiring access to PTs, but which might still provide a "projected formula" appropriate for examples like: (11) The cat walked round the toy lion.

\* those familiar with the system of Wilks (1968) 1965 etc.) will remember that these are the "bare template" structures actually used to obtain the initial template match. The suggestion here is that the "knowledge-aspect" of these highly-general structures is to be found as the pseudo-texts of primitives - as the latter function right at the top of the conceptual hierarchy imposed by the thesaurus.

Then he came back and sniffed it.  
Where we might be helped to refer "he" and "it" correctly by the new, projected, formula [lion] whose head was no longer BEAST, and which could therefore no longer be the reference of "he" as a real lion would be.

A more radical and interesting development would be the construction of "PT repacking functions" specific to certain qualifiers. Thus, for example, such a function for "toy", if faced with the phrase "toy car" might repack (7) using a general rule to delete all constituent templates based on the action USE, as well as all those that are at end of a GOAL tie, since toy cars cannot, normally serve human needs, uses and purposes.

### An Implementation Environment

It is proposed to implement these suggestions currently using a new base for the preference semantics system, expressed as a semantic augmented transition network. It is hoped that this more conventional re-expression of the semantics-driven pattern matching parser will make it more comprehensible and acceptable. The environment will be a question-answering system that discusses in English, the construction of the semantic representation itself for a text. It is hoped too that this will avoid the trivialities that lurk in answering questions about mini-worlds, where the answers to all possible questions sometimes seem preformed in the data base. Some simulation of parallel processes is hoped for, so as to distinguish (as independent matching processes) those aspects of the present system that correspond to local context (like preference) and those that correspond to a global context (like PTs and the use of the thesaurus generally), and to apply these independently to a text, as parallel branches of analysis, each seeking template-like structures in the input text. The interest, and the possibility of test of the relative usefulness of the branches, would come from switching off some particular branch and seeing if the analysis deteriorated.

Let us imagine such a dialogue as it might apply to the analysis of a real newspaper text like: (The Times: 5.2.76).

- (12) An ambulance driver, answering an emergency went through red traffic lights and collided with a car killing its driver. The interest here is how the program is to determine the reference of "its" as the car, rather than the ambulance driver. In what follows, the user interventions are in the upper case, notes are preceded and followed by square brackets.

END OF TEXT INPUT

1. An ambulance answering an emergency call
2. S went through red traffic lights
3. and S collided with a car
4. S killing its driver

[S marks missing template agents in the four fragments the text has been cut into for ease of reference]

All fragments match templates.  
Template 2 fails preferences for agent.  
Now repacking S slots

2. The ambulance driver went through red traffic lights.

WHY?

a driver is mobile, a call not.

3. The abulance driver collided with a car.

Suggest:

4. The ambulance driver killed ? its driver

[system unsure of S in (4), it is marked as a problem]

WHY?

only active agent available

But now re-inferring as follows:

preference breaking in (2) drives inferences as follows:

traffic lights are aperture

or

"went through" means "passed"

I select latter and project on to "went through"

[straightforward extraction in (2) plus access to pseudo-text for "traffic-lights"]

examining preferences in

4. S killing its driver.

"It" is lights (number wrong)

or

call

or

car

or

ambulance

Prefer car or ambulance for "it"

[calls don't have drivers]

Confirm

4. ambulance driver kills car's driver

WHY?

First:

from traffic-lights pseudo-text

passing red traffic light

causes people hurt

["pass" has replaced "went through" by projection, remember]

[pseudo texts can have a "cause tied" template pair read out]

this confirms kills car and ambulance driver.

Second:

Calling pragmatic over-ride

its=cars

"its" not twice coreferential

[a general Gricean\* style pragmatics rule that maximises information in the fragment.]

Confirm its=cars

Representation complete.

The new point brought out of the hat at the end, as it were, shows the need for general pragmatic principles at the highest level. Just as we need preference at the lowest. The general interest here would be the possibility of two clashing general principles in the analysis of a given text: preference seeking, in some sense, to minimise information (argued in Wilks 1975), and another seeking to maximise it.

#### Discussion

The cause-tied template pair, read out from the PT for fragment 4, shows that the function of the PT is not tied to just simple matching and inferences, and can span across texts of some length.

Moreover, in an actual implementation, the hierarchical organization of the PTs would plan a stronger role than appears here. As was suggested in connection with example (10), the PTs would not all be stored explicitly, but would be constructed as required \* from the more general PT for the corresponding thesaurus sub-head term, together with appropriate slot fillers. Thus, the "ambulance example" would always access the ~~#~~ vehicle PT initially and would attempt to deal with the text at that more general level, before proceeding to construct the "ambulance PT" from the ~~#~~ vehicle one.

A number of very general comparisons and issues suggest themselves here. The PTs are clearly of the more static type of frame adumbrated in (Minsky 1975), in that, unlike those of (Charniah 1975) (Schonk 1975b) etc., they do not have a narrative line, temporally or causally ordered. Indeed, (7) is essentially unordered, merely connected, and, should unconnected sub-PTs exist in a PT, they may be considered connected by an unordered AND predicate.

So, since PTs are non-narrative, or in crude terms, for nouns not verbs, the question of parsing text expectationally with their aid does not arise (as has been suggested in Schonk 1975b, and argued

against in Wilks 1977b), although this maybe a purely academic question in that scripts are at present (Riesback 1977) being applied independently at parsing procedures. This leads to close analogies between the present paper and the proposals of (Granger 1977) to use knowledge-based techniques (though not actually the scripts available in the system he is working as part of) to understand unknown words in texts. The relation between unknown words and preference-violating known words is complex and beyond the scope of this discussion.

Another and very general and relevant issue is that of the relation of the "understood structure" derived for a text, and the memory structure to be associated with it. Ortóny (1975) has argued persuasively that one must not just assume the two to be the same, in the face of much counter-evidence. In the present proposals, the "episode" structure for a text has the same format as a PT, or memory structure, but the two are not assumed to be the same (or even differently filled in copies) for any input.

The implementation environment, or debugging aid for NLUS's, is, like Moch AI, only hand-waving in the face of very difficult, and ill-understood problems. It has a main distinguishing feature, it is in the attempt to separate those projections where the highly specific PTs are helpful from those where they are not. In some cases general principles, like extraction or "pragmatic override", seem sufficient. The key interest of the proposed implementation would be the possibility of assessing the relative values of highly specific structures and general principles. This is in keeping with the "lazy system" assumption that has always been behind this NLUS\*\* that a system should do no more processing and inference work than is necessary to deal with the analysis problem in hand, even if that problem is coping with extended usage.

It is hoped that a system able to project in this way, from both general and specific knowledge structures, and to relate such projections to the application of wholly general pragmatic principles might give insight into the complex role of knowledge in language understanding, and into the peculiar role of language boundaries that the preference restrictions symbolize.

\* (Grice 1967) The principle of avoiding unnecessary redundancy: if "it=ambulance driver" the writer would have used "killed himself".

\*\* Cf. the notion of "variable processing depth" in Bobrow & Winograd.

# REFERENCES

- Robrow, D & Winograd, T. KRL - an overview of a knowledge representation language. Cognitive Science 1, 1977
- Givon, T. The structure of ellipsis. (Systems Development Corp.: Sta. Calif). 1967
- Granger, R. Foul-up.ubmitted to Fifth IJCAI, 1977
- Grice, H. Logice and Conversation, Unpublished mss. 1967
- King, M. & Wilks, Y. Semantics, Preference and Inference. (Inst. for Semantic and Cognitive Studies: Geneva). 1977
- Minsky, M. A framework for represeantion knowledge in Winston (ed.) The Psychology of Computer Vision. McGraw Hill: New York). 1975, 211-277.
- Ort6ny, A. How episodic is semantic memory? Proc. Theoret. Issues in Natural Language. Pricessing. BBN: Camb. Mass. 1975.
- Riesbeck, C. (as Granger) 1977
- Schank, R. (ed.) Conceptual Information Processing. (North Holland: Amsterdam). 1975a
- Using knowledge to understand (as Ort6ny). 1975b
- Wilks, Y. Computable Semantic Derivations. (as Givon). 1968
- A preferential, pattern-matching, semantics for natural language understanding. Artificial Intelligence, 6, 1975.
- De Minimis: the archaeology of frames. Proc. AISB Conference (Dept. of A.I., University of Edinburgh), 1976a.
- Processing Case. Amer. Jn l. Comput. Ling. 56. 1976b.
- Making Preferences more active, Edinburgh Dept. of A.I., memo No.32, 1977a.
- Wilks, Y. Frames, Scripts, Stories and Fantasies Pragmatics Microfiche, 1977b.