SHIFTING MEANING REPRESENTATIONS

Karen Sparck Jones

Computer Laboratory, University of Cambridge
Corn Exchange Street, Cambridge CB2 3QG, England

ABSTRACT

This paper argues that a number of different kinds of meaning representation, between which partial translations can be made as needed, are all required for a reasonably comprehensive language processing system. These representations capture different and possibly complementary aspects of a text's form, content and reference worlds, and are suited to different subtasks of the language processor. Initial testing of the proposition via a system designed for natural language access to databases is described, showing how different types of meaning representation with different characteristics are called for, related, and used.*

1. Introduction: the single representation assumption

It has been widely assumed in natural language processing research that a reading for a piece of text has a single meaning representation. When a cycle of morphological, syntactic, semantic and pragmatic processing has been carried out for some piece of text, the resulting representation embodying the system's interpretation, or 'understanding', of the text is the only form of representation proper to the text. A representation may be modified, for example if the processing of a later sentence supplies a specific referent for a definite noun phrase. But it is not changed in kind. Thus if the system's language-processing task depends on a long-life, large-scale text representation, this may be modified in the course of processing by the arrival of new information; but it is not replaced by a different type of representation.

This assumption underlies approaches working with representations of quite different content and form, and both approaches working with 'shallow' representations and those working with 'deep' ones, defining the former very broadly here for illustrative purposes as representations close to the actual text structure, and the latter as representations which do not attempt to preserve the given order, or individuality, of text items. The general approach so far has been to choose a means of representation which appears sufficiently informative, but abstract, to support the input-text-driven operations required to carry through the system's particular task. Thus for natural language translation, for example, relatively inexplicit and shallow representations have been deemed adequate, while for systems involving more extensive inference, for example question answering, more explicit and deeper representations are commonly used. Some systems involving language may be sufficiently restricted or specialised to fall outside this paradigm, by not posing serious requirements for text meaning representation: for example systems exploiting simple input texts to modify an independent database not intended to record linguistic features of the text.

*This work is supported by the U.K. Science and Engineering Research Council.

2. The multiple representation requirement

It is, however, at least possible that full text interpretation, even for the very modest task systems that are all we can tackle at present, requires more than one meaning representation. That is, the system requires representations which are all meaning representations because they are interpretable in an underlying world, but which are couched in different representation languages having distinctive characteristics appropriate to particular subtasks of the overall task. The relation between the representation languages may be quite complex, and distinct but related underlying worlds may be involved. The important point is that the different meaning representations all have genuine status within the system: we are not concerned with the case where several representations, which could all be deemed meaning representations, are constructed, but all but the final one are treated as disposable intermediates. I claim that when a language processing system's task requirements are investigated, they imply that the system must be able to manipulate, simultaneously and on demand, several different types of meaning representation for a text, conveying different and even exclusive kinds of information.

These representations cannot be mere notational variants of one another. If the task system is to make use of them for different subtasks, there must be procedures for moving, or translating, from one form to another; but this does not imply that all of the properties of the text expressed in either of a pair of representations between which translation is being made are to be genuinely found in the other. It is important here to distinguish proper, functional, elements of a representation from mere baggage. For example, one representation may preserve specific lexical items, the other indicate only underlying primitive concepts; attaching lexical items to the latter to ensure some recovery of the former does not necessarily imply that the lexical items in the latter have any true representational function. The assumption is that for any pair of meaning representations, partial, but not complete, equivalence exists, permitting transitions between them. Saying that this set of representations nevertheless collectively constitutes a single meaning representation is vacuous.

The motive for the view that more than one representation is required is supplied, on the one hand, by tasks like machine translation. It is evident that natural language translation (at least for some purposes) must seek to preserve text structure and expression, i.e. text narrative form. It is equally evident that correct translation can require arbitrarily extensive inference, for example to identify pronoun referents *or* compound noun structures, and quite possibly both 'weak' inference of the kind illustrated by Wilks and 'strong' inference dependent on explicit quantification. The complementary motivation is supplied by observations like Partee's, that an explicit quantification of a sentence like "More people get more satisfaction out of L&M's than out of any other cigarette" is so complex that it doubtfully reflects the simple, if furry, meaning accepted by the reader.

Taking predicate-logic-type representations as deep, examples like this suggest that the different meaning representations may have different depths, and specifically, that there is a single sequence of representations. But cases like translation suggest that different representations at similar depths are also needed, and specifically shallow representations for the input and output natural languages, which are directly related to one another as well as to deeper representations. Thus it is more important to focus on differences of representation content than of depth.

3. Investigating multiple requirements

Our current project (Boguraev and Sparck Jones, 1983)t for database access, provides an appropriate environment for an initial study of the need for several meaning representations, and of mechanisms for moving from one to another. This follows from the approach to database access being adopted. The objective of a natural-language front-end processor for a database management system is to convert a natural language question into a (formal) search query appropriate to the particular data model used for the database: e.g. into an expression in relational algebra. The normal processing strategy adopted for this purpose involves applying a data-domain-specific semantic processor either alongside, or largely subsuming, conventional syntax. The present project's aim is to avoid having to provide a new semantic processor operating directly on input text for every new database. The approach adopted assumes that a good deal of semantic *as* well as syntactic processing can be achieved using domain-independent information, and that the domain-specific processing which is necessarily required for the database access task is reduced in scope and simplified by being able to exploit the output of the earlier general syntactic and semantic processing of the question.

In the system which has been constructed, the initial general interpretation of the question provides a rich meaning representation of the input as a case-labelled dependency tree with slots filled by word senses defined by semantic primitives. This is then processed, still as a general, database-independent operation, to derive a predicate-logic-type representation in the LUNAR style (Woods, 1978), retaining the case relation and category primitive information, but discarding linguistic information like clause boundary markers. The logic representation is in turn processed using domain-specific information to map the case and category information onto the object, property and relationship concepts of the database world. The formal query representation thus generated is in turn converted into the actual search representation of the input which is a relational algebra expression geared to the administrative organisation of the database. All these operations are semantics driven, but are inevitably different in their essential character, just because they are designed to build different representations, and are too complex to detail here. (They have already been quite extensively tested on a sample database.)

In the present context, the relevant properties of this natural language processor are first, that the various representations of the input question are all meaning representations. Because the initial representations for the input have a perfectly good direct representation in the ordinary world, via the case and category primitives, it cannot be said that the only meaning representation for the question is the final search query. The approach of LUNAR and PHLIQA1 (Bronnenberg et al, 1979) is not being followed here. Equally, both query and search representations *refer* to the database world. Second, it is not the case that the four representations differ only in the domain independence or dependence of their semantics. This in itself constitutes an important difference of representation content; but there are also significant differences of content and form between the initial question representation and the logic representation, which exploit different view of the ordinary world, and between query and search representation, which exploit different views of the database world. These differences may also be seen as differences of depth: the initial representation is in an obvious sense shallower than the logic representation; but this is not the difference of real importance here.

The fact that the processor generates just four explicit representations is arbitrary: this is a product of experimental system design. No claim is being made for the individual details of the various representational forms. What matters is that the task, which is an instance of a widespread language-using activity, namely inquiry, presupposes references to different worlds - the general and the special, and uses different forms of representation relating to each of these worlds - (loosely) the linguistic and the logical. To argue that we are really operating in one world, constructed from two source worlds, is a formalist view which fails to capture the fact that it is the construction process itself which is embedded in the task; more importantly, the representations are not notational variants, as information is thrown away in the transitions from one to another.

That so far, these transitions are made in one direction only, is irrelevant. While the first requirement of a database access system is for a straightforward search equivalent of the input natural language question, it is obvious that a fully-developed access system will involve other interactions between representations, and transitions, in other directions. These will be required, in particular, first to meet the inference needs of question interpretation, and second, to respond effectively to the user.

Inference processes have not yet been designed for the system, and it is not clear in detail what form they should take. However a beginning has been made on response processes. One obvious response is checking the question interpretation with the user. For this, reformulation of the question reflecting each stage in its processing has its own justification. For example, text paraphrases generated from the first, dependency representation can show the user, in language close to his own, how linguistic ambiguities have been resolved; generation from the logic representation can check the desired scope of quantification, and generation from the query and search representations the way the user's question is being viewed in the database context. Figure 1 shows examples of texts generated first from the initial dependency representation, and second from the final search representation, for a database question. The first indicates disambiguation, the second makes obvious reference to the relational structure of the database. The text differences reflect the different functional characteristics of the underlying source representations. Generation here is done directly; it does not have to proceed through inverses of the interpretation and where this might be appropriate, it would not imply that individual input representations were recovered: only ones as near as the representation language relationships allow.

In the project work so far, only translation in the most straightforward direction from one representation to another has been implemented. Jut in general, it appears that interpreting, in the full sense of establishing the legitimacy of, an input question, not to mention answering it, can involve all types of representation, and hence reruire free movement from one to another.

4. Conclusion

Investigation of other language processing tasks, like summarising, supports the argument for multiple representations: far-ranging and heterogeneous inference operations may be required to identify important parts of a text, but there is also a need, for at least some abstracting purposes, to retain global narrative structure. The essential language processing model being proposed here is therefore one using a set of meaning representations for an input (or output) exhibiting distinct, and in the limit exclusive, properties of the source text, with translation between these exploiting different specific common types of element for particular transitions, as required to meet various interpretive subtasks. The present project aims to show what different representations are needed and what mechanisms are required, working collaboratively, to use them.

The idea of multiple representations is in some ways similar in motivation to the idea of multiple descriptive viewpoints embodied in the design of KRL (Bobrow and Winograd, 1977). But Bobrow and Winograd do not consider the systematic exploitation of these ideas for the purposes of producing different types of non-equivalent text meaning representation as considered here, or the procedures required to move from one to the other.

ACKNOWLEDGEMENT

I am grateful to John Tait for discussions.

REFERENCES

[1] Boguraev, B.K. and K. Sparck Jones, "How to drive a database front end using general semantic information." In <u>Conference on Applied Natural Language Processing.</u> Santa Monica.California, 1983, PP81-88.

[2]. Woods, W.A. "Semantics and quantification in natural language answering." <u>Advances in Computers</u> 17 (1978) 1-87.

[3]. Bronnenberg, W.J.H.J. et al. "The question answering system PHLIQA1." In <u>Natural language question answering systems</u> (Ed Bole). London: Macmillan, 1979-

[4] Bobrow, D.G. and T. Winograd, "An overview of KRL, a knowledge representation language." <u>Cognitive Science</u> 1 (1977) 3-^6.

"Get all status 30 red part suppliers."

(1) "Find all red part suppliers with status 30."

(2) "Show me the key of the suppliers whose status is 30, and who make shipments of parts whose colour is red."

FIGURE 1