

NLP: SEMANTICS AND KNOWLEDGE REPRESENTATION

Lenhart K. Schubert
Department of Computer Science
University of Rochester

In a N[atural] L[anguage] P[rocessing] context, the terms ‘semantics’ and ‘semantic interpretation’ usually refer to methods of representing the meanings of natural language expressions, and of computing such meaning representations. This brand of ‘syntactic semantics’ is distinct from the notion of denotational semantics as understood in formal logic. (This is not to say that the latter notion is irrelevant – see below.)

Semantic interpretation in many NLP systems can be viewed as involving three interleaved phases: determination of a preliminary, underspecified semantic representation of the input, disambiguation of this representation, and inference of implicit or ‘missing’ information, notably connections between causes and effects, means and ends (goals), and between parts and wholes (things, episodes, plans). Other kinds of implicit information which need to be extracted for genuine discourse understanding are implicatures, presuppositions, and inferences about the speaker’s or author’s knowledge, goals, and plans, and about the structure of the discourse and context. [*See Discourse, Pragmatics, Implicature, and Presupposition.*]

In SYNTAX-DIRECTED (or syntax-driven) systems, the preliminary representation is derived from a syntactic analysis of the input, using rules that relate components of the syntactic analysis to components of the meaning representation. In SEMANTICALLY DIRECTED systems, text is mapped directly into parts of meaning representations, often using MEANING TEMPLATES associated with particular words, and procedures which seek fillers for the parts of those templates.

DISAMBIGUATION includes choosing among alternative preliminary representations (corresponding to alternative syntactic analyses), selecting word senses, scoping quantifiers, resolving anaphora, and choosing among alternative ways of filling in implicit connections.

The recognition that inference of implicit connections plays a central role in language understanding is one of the main insights that has emerged from NLP research. Understanding is thought to be achieved by processes that match input information against stored knowledge (axioms, rules, plans, ‘frames’, ‘scripts’, etc. – see below), allowing the new information to be elaborated and linked with previously given or derived information. It remains unclear what general principles guide these processes; one suggestion is that the overall goal is to find the ‘least costly’ (or most probable) explanation of what has been said (or why it was said) (see Hobbs *et al.* 1993). Experience in NLP indicates that the requisite body of knowledge – about word meanings, about discourse conventions, and most of all, about the ‘world’ (the domain of discourse) – is very large for non-trivial domains.

Besides the distinction between syntax-directed and semantically directed systems, another major distinction among interpretive strategies concerns the way in which they treat semantic coherency constraints such as selectional restrictions. They may be treated as axiomatic knowledge, as constraints on the syntax-semantics mapping, or as constraints on syntactic form (as in so-called semantic grammars). A good general source of information on semantic interpretation is Allen 1995, parts II and III.

1. The form of semantic representation. The above sketch of the semantic interpretation process leaves open the question of what form the semantic representation of an input, and the knowledge needed to compute it, should take. Semantic representation and knowledge representation have generally been assumed to require a well-defined symbolism along with a theory of how this symbolism can express information about a domain of discourse. However, a general distinction is made between DECLARATIVE representations and PROCEDURAL ones. Declarative representations are thought of as encoding propositions (facts, beliefs, etc.), while procedural representations encode potential behaviors (methods, skills, techniques, etc.). Most frameworks for knowledge representation allow for both types of representation, though they may bias implementation choices one way or the other.

In discussions of declarative representations two major controversies repeatedly surface. One is whether such a representation should admit a Tarski-style denotational semantics and theory of truth. ‘Logicians’ insist that it should, if it is to be comprehensible and if it is to support theories of sound inference; their opponents question the feasibility and usefulness of the logicist enterprise, and seek to justify representational choices either in terms of procedural efficacy or in terms of cognitive and linguistic consider-

ation (see for example Jackendoff 1990 for the latter perspective).

The other controversial issue is whether word meanings should be decomposed into ‘primitive’ elements in the process of mapping natural language input into an internal representation. It is argued that decomposition is required in order to capture shared inferences, such as that both eating and drinking result in a transfer of material to the inside of the agent (as may be inferred from decompositions based on an *INGEST* primitive). A counterargument is that such generalizations can just as well be captured by axioms relating specific predicates (such as *EAT* and *DRINK*) to more general ones (such as *INGEST*). This avoids the problematic assumption that word meanings are definable in terms of a relatively small number of primitives.

2. Common Types of Representation Perhaps the simplest declarative representations are those based on sets of *ATTRIBUTE-VALUE* pairs attached to objects (words, things, facts, events). Several early question-answering systems used this type of representation, and many of the more advanced types of declarative representation can be viewed as elaborations of attribute-value representations.

RELATIONAL DATABASES, from the present perspective, are essentially large-scale attribute-value tables. They have been the focus of much activity in NLP, because of the attractiveness of natural-language access to information in existing databases, and the relative ease of interpreting and answering questions confined to database contents. Moreover, the importation of AI techniques into database systems, and vice versa, has led to a convergence toward a common ‘knowledge base’ area (see for example the articles in von Luck and Marburger 1994).

LOGICAL REPRESENTATIONS, whose origins lie in the traditions of philosophical and mathematical logic, were proposed early as both semantic representations and as representations for commonsense knowledge, including knowledge about actions, abilities, goals, and causes. Early work on NLP tended to view the mapping from the syntactic form of sentences to their underlying logical form as rather haphazard and complex. Several later developments shed new light on this mapping and facilitated computational realizations. Montague Grammar [q.v.] illuminated intensional locutions (such as *seeks a unicorn*) and showed how to derive logical forms from syntactic forms compositionally, in node-by-node fashion. The combination of Montague’s ideas with new grammatical frameworks such as Generalized Phrase Structure Grammar [q.v.], Head-driven Phrase Structure Grammar [q.v.] and Categorical Grammar [q.v.] proved particularly fruitful for computational linguists. Situation Semantics [q.v.] has clarified the relation

between sentences and the situations they describe, along the way providing an account of how logically equivalent sentences can differ in meaning; various versions of its metalanguage for specifying and relating situations have been developed into meaning representations for use in NLP work. Discourse Representation Theory (DRT) [q.v.] tackles other long-standing problems, including that of providing a formal account of anaphoric dependencies in multi-sentence texts, and more generally that of explicating the way in which utterances depend on and add to the meaning of the prior discourse. DRT has also given rise to many implementations. There is also a clear trend towards amalgamation of ideas from various approaches to representation, including those just mentioned, into more comprehensive frameworks and computational systems. Logic programming has also facilitated the syntax-semantics transduction, by allowing all types of knowledge needed for interpretation – about syntax, semantics, pragmatics, and the world – to be logically encoded.

Like logical representations, PROCEDURAL REPRESENTATIONS have some philosophical precedent, particularly in logical positivism and Ludwig Wittgenstein's view of language as a kind of communicative game. Semantic interpretation of an input, on the procedural view, involves construction of an executable command such as an information storage or retrieval request. Many database-oriented systems have interpreted inputs in this way, including LUNAR (Woods 1977) and its many descendants. LUNAR also encoded grammatical knowledge procedurally; and the 'blocks world' program SHRDLU (Winograd 1972) took the further step of proceduralizing world knowledge. Thus, the fact that all (toy) blocks are manipulable was recast imperatively as 'If you wish to show that x is manipulable, try showing that x is a block', or alternatively as 'Upon learning that x is a block, conclude that x is manipulable'. A more recent development is the attempt to use schematic descriptions of sensory-motor processes not only as a means of specifying an agent's behavior but also as a basis for interpreting language (Narayanan 1997).

SEMANTIC NETS originated in psychologically oriented studies of language understanding, but have also found application in many practically oriented projects (e.g., see Lehmann 1992). They are characterized by (i) a graph-theoretic propositional syntax closely related to logical syntax (though perhaps motivated by cognitive or linguistic hypotheses), (ii) a PROPERTY STRUCTURE organization which makes propositions about a thing directly accessible from a token for that thing, and (iii) processes such as SPREADING ACTIVATION from referenced nodes (as a basis for semantic disambiguation)

and PROPERTY INHERITANCE in generalization hierarchies. It should be mentioned that some form of (ii) and (iii) is found in most natural language understanding systems; furthermore, whether or not a computerized propositional representation is seen as graph-theoretic is to some extent in the eye of the beholder.

FRAMES are structures encoding knowledge about stereotyped kinds of objects or situations, with special provision for the ‘roles’ played by their parts or participants (Minsky 1975). Frames used in semantic representation are often termed CASE FRAMES, reflecting the influence of theories of CASE RELATIONS [*See Case.*] When used in conjunction with semantically directed interpreters (as is common), case frames are instantiated by procedures that seek semantically appropriate fillers for the case roles. Frames are also in common use for world knowledge representation, especially for specifying hierarchies of concepts and their ‘roles’, with semantic net-like property inheritance. As such, their usefulness lies in the systematic way in which they encode general expectations about the attributes possessed by concepts, including the default values of those attributes. Closely related to frames, in this respect, are SCRIPTS and PLANS, structures that provide general expectations about the way familiar kinds of events typically unfold, and familiar kinds of goals are typically pursued (Schank and Abelson 1977).

Finally, specialized representations for taxonomic, spatial, temporal, numeric, and other pervasive types of relationships appear to be indispensable for building NLP systems with human-like inferential capabilities. Such representations can be viewed as analogs, in certain crucial respects, of the relations they model. The practical advantages of such representations have long been recognized, and there appears to be evidence for their psychological plausibility as well (e.g., Johnson-Laird 1983).

3. Research Issues Most of the topics touched upon here remain under intensive investigation at the time of writing. As far as methodological issues are concerned – and these include the viability of logicism, of decompositional approaches to word meaning, and of various declarative and procedural formalisms – no resolution or consensus can be expected on any theoretical grounds. Experience with more powerful NLP systems will be the final arbiter.

In the meantime, all proposals for semantic representation still confront many specific difficulties, such as the representation of tense and aspect, of adjectival and adverbial modification, nominalization, generic sentences, propositional attitudes, counterfactual conditionals, comparatives, and generalized quantifiers. Consequently, rules for obtaining preliminary semantic

representations of natural texts can be given only for very restricted subsets of natural languages. As a result, much practically oriented research has focused on the development of techniques for extracting very specific information items from text corpora by pattern matching techniques that do not depend on deriving a full semantic representation, or on making inferences [see Applications of Natural Language Processing]. Similarly, many aspects of the disambiguation process remain obscure, though statistical approaches to this problem have made significant headway in recent years on both structural and word-sense disambiguation [see Word Sense Disambiguation]. The inference of implicit causal connections, plans, goals, reasons, and so on remains a refractory problem.

The number of semantic and factual details underlying interpretation and understanding is so huge that much current research looks towards LEARNING, often based on scanning large computerized corpora of annotated or unannotated texts, as a way of bootstrapping NLP systems towards levels of increased competence.

LENHART K. SCHUBERT

Bibliography

- ALLEN, JAMES. 1995. *Natural Language Understanding*. Redwood City, CA: Benjamin/Cummings.
- DALRYMPLE, MARY, JOHN LAMPING, FERNANDO C.N. PEREIRA, and VIJAY SARASWAT. 1995. Linear logic for meaning assembly. *Workshop on Computational Logic for Natural Language Processing (CLNLP-95)*, Apr. 3-5, Edinburgh.
- HOBBS, JERRY R., MARK E. STICKEL, DOUGLAS E. APPELT, and PAUL MARTIN. 1993. Interpretation as abduction. *Artificial Intelligence* 63:69–142.
- JACKENDOFF, RAY. 1990. *Language and Cognition*. Cambridge, MA: MIT Press.
- JOHNSON-LAIRD, PHILLIP N. 1983. *Mental Models*. Cambridge, MA: Harvard Univ. Press.
- LEHMANN, FRITZ (ed.). 1992. *Semantic Networks in Artificial Intelligence*. Oxford: Pergamon Press.

- MINSKY, MARVIN. 1975. A framework for representing knowledge. In *The Psychology of Computer Vision*, edited by Patrick H. Winston, pp. 211-280. New York: McGraw-Hill.
- NARAYANAN, SRINIVASAN. 1997. Knowledge-based Action Representation for Metaphor and Aspect (KARMA). Ph.D. thesis, University of California, Berkeley.
- SCHANK, ROGER C., and ROBERT P. ABELSON. 1977. *Scripts, Plans, Goals and Understanding: An Inquiry into Human Knowledge Structures*. Hillsdale, NJ: Lawrence Erlbaum.
- VON LUCK, KAI, and HEINZ MARBURGER (Eds.). 1994. Management and Processing of Complex Data Structures. *Third Workshop on Information Systems and Artificial Intelligence*, Hamburg, Germany, February 28 - March 2, Proceedings. *Lecture Notes in Computer Science*, Vol. 777. New York: Springer.
- WINOGRAD, TERRY. 1972. *Understanding Natural Language*. New York: Academic Press.
- WOODS, WILLIAM. 1977. Lunar rocks in English: explorations in natural language question answering. In *Linguistic Structures Processing* (Fundamental studies in computer science, 5), edited by A. Zampolli, pp. 521-569, Amsterdam: North-Holland.