

Syntax-Driven and Ontology-Driven Lexical Semantics

Sergei Nirenburg and Lori Levin

*Center for Machine Translation
School of Computer Science
Carnegie Mellon University
{sergei, lsl}@nl.cs.cmu.edu*

Abstract

In this position paper we describe the scopes of two schools in lexical semantics, which we call syntax-driven lexical semantics and ontology-driven lexical semantics, respectively. Both approaches are used in various applications at The Center for Machine Translation. We believe that a comparative analysis of these positions and clarification of claims and coverage is essential for the field as a whole.

There are different traditions in the study of lexical semantics. Two of them seem to be the most current in computational linguistics and its applications — the one based on syntactic theory studies and the other, on the artificial intelligence approaches. The former seeks to discover semantic properties of lexical items from which syntactic behavior (such as subcategorization and participation in transitivity alternations) is predictable. (See Grimshaw, 1990; B. Levin & Rappaport Hovav, 1990.) The latter tries to establish the meaning of natural language texts with the help of an independently constructed “world model” (often called “ontology”) which explicates relations among entities in the world rather than lexical units.

It is customary to pitch the two approaches as competing. And in practice, researchers typically develop lexical-semantic theories and computer systems based on them using only one of the two. We would like to argue, however, that these approaches to semantics are much closer to one another in their aims and research methodologies than to any other schools of semantics (“logical,” “philosophical,” “formal,” etc.) From a practical standpoint, we also have a different experience to report. We believe that neither on theoretical grounds nor with respect to computational applications is there a necessity to make one of the approaches prominent at the expense of the other. We have carried out both theoretical and practical work in which the approaches seem to coexist. The central role of the syntax-driven lexical semantics in the process of NLP is to decode the nature of the dependency between heads of phrases and their arguments in a particular language. This knowledge is then used in ontology-driven lexical semantics as a necessary set of heuristics which allow us to represent the meaning of a text in terms of a language-independent conceptual model. Thus, we believe that a comprehensive approach should combine the benefits of both approaches and that neither will on its own be sufficient for realistic NLP.

1 Syntax-Driven Lexical Semantics

In a knowledge-based system, the primary role of syntactic analysis is to decode knowledge that is syntactically encoded. An important component of such knowledge is the

relationship between predicates and their arguments. For example, in order to know that the English sentence *Max interviewed Hester for a job* means that Max was considering hiring Hester and not that Hester was in the position to hire Max, it is necessary to know that the interviewer role is expressed as the subject of the sentence, that the interviewee role is expressed as the object, that the subject precedes the verb and that the object follows it. Much of the information about the syntactic encoding of arguments can be placed in lexical entries of argument-taking predicates. Classes of predicates with similar argument encodings usually have some semantic similarity, and thus define a kind of lexical semantics based on semantic features that are predictors of syntactic behavior.

In light of these observations, our theory of lexical semantics includes a component of lexical knowledge that describes the syntactic encoding of arguments. This lexical knowledge consists of Lexical Conceptual Structures (LCS) and Linking Rules. We will assume that the LCSs of any given language are composed from a universal set of building blocks (including, but not limited to those described by Jackendoff, 1983, 1990). However, the presence of lexical gaps—meanings that are lexicalized in some languages but must be expressed phrasally in another language—suggests that languages do not all share exactly the same inventory of LCSs. The most important property of the LCSs of any individual language is that they contain all of the information necessary for Linking Rules to calculate the syntactic encoding of predicates and arguments. Thus they contain syntactically relevant semantic information.

Linking Rules describe how LCSs are syntactically encoded in a language and they partition LCSs into classes of predicates that have similar syntactic encodings. Linking Rules and the verb classes they define are language-specific (Mitamura, 1989), but governed by language-universal principles. The particular version of linking rules and verb classes we have implemented is based on Lexical Functional Grammar (L. Levin, 1987; Bresnan & Kanerva, 1989), and includes the treatment of the phenomena briefly discussed in the rest of this section.

1.1 Which Grammatical Function Is Associated with Each Argument?

Linking Rules associate grammatical functions with arguments in a LCS. For example, the English verb *interview* associates its interviewer role with the grammatical function SUBJECT and its interviewee role with the grammatical function OBJECT. Argument-taking predicates themselves fall into subcategories (or subcategorization classes) based on the grammatical functions that they assign to their arguments. For example, those that assign SUBJECT and OBJECT, those that assign SUBJECT and OBLIQUE, etc.

1.2 Which Syntactic Position, Case Marker, or Agreement Morpheme Encodes Each Argument?

After linking LCS arguments to grammatical functions, the next step in the syntactic encoding of LCS arguments is to assign a syntactic position or morphological marker to each argument. For example, in English, the interviewer role is syntactically encoded to the left of the verb. In another language, the verb might specify a case marking for the interviewer role. Default information about syntactic encoding of arguments (e.g., SUBJECTs have nominative case) can be specified in the syntactic rules of the language, instead of in the lexicon. However, after these defaults are specified, there remains a

considerable amount of lexically specific information—which prepositions mark oblique arguments, which verbs take dative objects, and so on.

1.3 Which Predicates Participate in Alternate Assignments of Grammatical Functions?

Many predicates allow multiple assignments of grammatical functions or case markings to their arguments. For example, the well-known *spray/load* alternation involves the assignment of the grammatical function OBJECT or OBLIQUE to an argument that is filled or covered, as in *spray the wall with paint* or *spray paint on the wall*. The passive construction also results from an alternation in grammatical function assignment to arguments. An argument that has the OBJECT function in an active sentence has the SUBJECT function in a passive sentence.

Argument-taking predicates fall into classes according to which alternations they allow. When a set of LCSs is identified that all allow the same alternations, it is usually possible to identify a semantic feature that they have in common that correlates with the transitivity alternation. Hence LCSs that undergo the same transitivity alternations appear to form semantic classes — for example, the change of state verbs, the spray-load verbs, verbs of removing, verbs of creation, etc. (See B. Levin (1989) for an extensive list of such classes.)

1.4 In Which Argument-Structure-Building Operations Does the Predicate Participate?

Talmy (1975, 1985), Jackendoff (1990), and B. Levin & Rapoport (1988) among others have proposed rules that conflate two lexical meanings to create a new lexical meaning. Such argument-structure-building rules in English include incorporation of manner of motion and directed motion (the argument structure for *jump* and the argument structure for *go* combine to form the argument structure for *jump into the room*); resultative secondary predication (the argument structure for *become* and the argument structure for *hammer* combine to form the argument structure for *hammer the nail flat*); and many others. For these operations, we are assuming that two LCSs are combined to form a new LCS, and this new argument structure undergoes assignment of grammatical functions and syntactic encoding of arguments. The allowable types of LCS conflation and the classes of verbs that undergo them vary from language to language.

All of the operations mentioned in this section identify syntactic classes of verbs—the verbs that undergo the same transitivity alternations, the verbs that map their arguments onto the same grammatical functions, the verbs that undergo a certain argument structure building operation, the verbs that encode their arguments syntactically in a certain way. It is usually possible to find semantic features of LCSs that correlate with these classes—e.g., a change of state, an effect on a patient, etc. In this sense syntactic patterns of verb classes define a lexical semantics. This lexical semantics is language-specific in that the membership of verb classes and the semantic features that they are based on are different in different languages.

2 Ontology-Driven Lexical Semantics

While the focus of attention in syntax-driven lexical semantics is on components of lexical meaning that determine syntactic behavior, the task of ontology-driven lexical semantics is to determine and specify meanings of lexical units, suggest how they must be represented (in the lexicon) and how they contribute to the overall representation of text meaning. We view this task as a component of the more general task of representing the meaning of a text in a computational application. This view has been developed mostly in the tradition of natural language processing within artificial intelligence. (It is impossible to reference all of the contributors to this, by no means monolithic, paradigm; a very incomplete list of references might include Wilks (1975), Schank (1973), Sowa (1984), Hirst (1987), articles in Hobbs and Moore (1985) or Lehnert and Ringle (1982)).

Some of the desiderata for an ontology-driven lexical semantics are discussed in the remainder of this section.

2.1 It Must Be Constructive

Among our differences from a typical practice in model-theoretical semantics is our “constructive” attitude to the models. We believe that in order to be able to discuss model-theoretic semantics, the models (that is, the ontologies) must first be constructed and put in correspondence with the lexicon. This premise also sets us apart from some work in linguistic semantics, for instance, Jackendoff’s. At the same time, we do not strive to use a very limited set of primitive concepts.

One way of constructing a model is to come up with a set of properties describing things in the world, define value sets for these properties and then describe each concept in the world as a set of particular property-value pairs. For reference purposes such sets can be named. The result is a world model, or computational ontology. A number of ontologies has been recently suggested and built, though not all of them with language processing in mind (cf. Lenat and Guha, 1990, Dahlgren and McDowell, 1989, etc.).

If some of the properties reflect links among ontological units, then the ontology can be viewed as a multiply interconnected network. Such a network can support “reasoning” about the world encoded in it, if special heuristic rules for network traversal (including inheritance) are defined (see, e.g., Woods (1975) Brachman (1983), Horty et al. (1990)).

2.2 It Must Support Disambiguation

In a computational environment, syntax-driven lexical semantics helps *distinguish* meanings but not represent them. Indeed, in some cases knowledge of predicate-argument structure can disambiguate a lexical unit entirely. Thus, the three senses of *spray* in LDOCE (illustrated by *spray paint on a wall*; *the water sprayed out over the garden* and *spray a wall with paint*) can be completely disambiguated on the basis of differences in their predicate-argument structure. However, *after* the choice has been made, the correct meaning needs to actually be represented. That is, we need a representation of meaning that shows *how* the different senses of *spray* are different, not simply that they are different.

However, in some cases a lexical unit cannot be completely disambiguated after the predicate-argument analysis is performed. Consider the English verb *hack*. We can distinguish two senses of the word — the cutting and the programming one. The difference

between these senses cannot be expressed in terms of the predicate argument structure. We will need to specify, for instance, that in one of the senses the instrument of the event is a knife or an axe and in the other, a computer. This type of information is, in fact, a familiar selectional restriction. Specifying selectional restrictions is different from specifying predicate-argument relationships. The specification of selectional restriction falls into the purview of ontology-driven lexical semantics.

In the lexis realm, ontology-driven lexical semantics provides the residue of meaning beyond the predicate-argument structure. Additional disambiguating power is to be provided on top of the verb class distinctions of the syntax-driven lexical semantics— for instance, selectional restrictions have to be formulated. Therefore, the lexicon should contain knowledge that is not required for decoding the syntactic realization of arguments. Naturally, this calls for additional expressive power in the formalism and, crucially, for reliance on an independently developed model of the world as the source of meaning interpretation for lexical units (see next item).

2.3 Metalanguage Must Be Different from Object Language

No matter how extensive a computational dictionary is prepared for an application, there will always be cases when a given set of selectional restrictions will fail to help disambiguate a text. This means that even more disambiguating power should be added to the arsenal of a lexical-semantic theory. The cases where selectional restrictions are not powerful enough are typically cases of metonymy or metaphor. Note that classification of phenomena as treatable by selectional restrictions or only by metonymy/metaphor processors crucially depends on the content of the dictionary — the less selectional restriction information in the dictionary and the fewer word senses distinguished, the more work remains for the metonymy/metaphor processor.

Since processing metaphor and metonymy is more effort-consuming than using selectional restrictions, it is natural that lexical semanticists have been searching for ways of reducing the need for the former. One way is to require very large and detailed lexicons. In part, this is the reason why some researchers suggested that word senses from a natural language could be used as elements of a metalanguage.

We believe in the distinction of metalanguage and object language and therefore claim that the meanings of natural language units should not be represented using unadorned lexical units of the same language. In the most general case, the meaning-encoding metalanguage should be a full-blown language-independent conceptual world model.

Constraints like selectional restrictions and distinctions among word senses have to be expressed using symbols — either words of a natural language or elements of an artificial “language of thought” (in the sense of Fodor, 1975). One problem with using words of a natural language is that they are typically themselves ambiguous and therefore a computer program which has to use them as disambiguation heuristics will run into additional disambiguation requirements.

Another reason for not basing an ontology on word senses of a natural language is that there are cases where different word senses, as specified in a human-oriented dictionary, seem to be divided quite artificially. Consider, for instance, the entry for *preempt* in LDOCE. It has four senses all of which conform to a single predicate-argument structure. All of them have something to do with having authority for preferential treatment with respect to goods or services. In fact, if one is so bent, one can think about distinguishing *more* than four such senses, at about the same level of generality. We believe that such

cases are prime candidates for merging the senses into one and treating the semantic analysis process for them in the metonymy/metaphor mode.

The above seems to argue that if automatic seeding of a computational lexicon is undertaken from an MRD, a nontrivial manual pass over the resulting lexicon should be carried out to make sure that the grain size of sense distinction is commensurate with the lexical and conceptual detail in the corresponding domain model.

2.4 Procedures for Assignment of Meaning Are an Integral Part of the Theory

Ontology-driven lexical semantics defines the meaning of a lexical unit through a mapping between a word sense and an ontological concept (or a part thereof). The process of text analysis, then, consists essentially of instantiating the appropriate ontological concepts or modifying property values in concepts that have already been instantiated.¹

2.5 It Must Offer a Treatment of the Phenomena of Synonymy, Hyponymy and Antonymy

In a theory where lexical meaning is expressed in terms of mappings into instances of ontological concepts, these standard lexical-semantic relations are explained in terms of relations among elements of a world model. Thus, synonymy will be defined as a relation among the lexical units at least some of whose senses map into the same ontological concept.² Lexical unit A is defined to be a hyponym of lexical unit B if the concept in which at least some of the senses of A map stands in the taxonomical *is-a* relation to the concept into which lexical unit B is mapped. Antonyms are defined on lexical units which map into regions on an attribute scale. Specifically, antonyms map into regions that are symmetrically positioned around the center point of the scale.³

2.6 Text Meaning Representation Language Should Not Be Too Close to the Syntax and Lexis of Any Natural Language

Since the purpose of an LCS in syntax-driven lexical semantics is to reflect the syntactically relevant semantic properties, that is those properties that determine how an LCS is mapped into a syntactic dependency representation, the LCSs are usually closely tied to the syntax and lexis of a specific language. They are therefore not as useful in multilingual computational environments as in monolingual ones. This is because languages differ in their lexicalization patterns and in patterns of syntactic dependency. Following are some examples that call for departure from the lexical semantics of individual languages in multilingual processing environments.

¹Lexical-semantic information is, of course, only one of several components of a text meaning representation. In practice, additional information (e.g., pragmatic and discourse-related) is encoded in lexicon entries and is represented as part of text meaning.

²It should be noted that in our application of lexical semantics, the DIONYSUS project, such mappings are allowed to be constrained, not just univocal (for instance, the English verb *taxi* maps, in one of its senses, to the concept *move-on-surface* but only if the theme of this instance is constrained to the concept of *airplane* or one of its ontological descendents) — see Onyshkevich and Nirenburg (1991) for additional detail.

³Note that synonymy on scalars is defined as mapping onto the same region of a scale.

- A head–modifier relationship can be realized in opposing ways in different languages; thus, a morpheme carrying the aspectual meaning can be realized as the head of a phrase (like the verb *continue* or *finish*), with a word denoting an event as its complement; alternatively, the word denoting the event can be realized as the main verb while the aspectual meaning can be realized through an adverb or even a affix.
- Lexical Gaps: Russian has no word that corresponds exactly to the English word *afford* (as in *I can't afford X* or *I can't afford to Y*). In a multilingual processing environment, there might be a concept corresponding to a sense of the English word *afford*. A Russian sentence *Ja ne mogu sebe etogo pozvolit'* (*I can't allow myself this*), uttered in a context of acquisition — which could have been assigned a straightforward lexical-semantic representation if we were building a lexical semantics for Russian alone — should involve the concept that represents *afford*. This means that if the units of the representation language are chosen so that they are based on Russian lexis, the meaning of *afford* will be missing. But this meaning seems sufficiently basic to be included in an ontology. As a result, if lexical patterns of many languages are used as the clues for building ontology, the quality of the latter should increase.
- Different languages describe the same event using different lexical semantics. For example, the ritual act of washing one's body in Islam is expressed with a change of state verb in Arabic (to get washed), as a verb of receiving in Indonesian (to receive ablutions), and with an agentive verb in French (to make ablutions) (Mahmoud, 1989).

For reasons such as the above, we argue that, at least in some applications, such as multilingual MT, a more language-independent representation scheme is preferable.

3 Our Position on Some of the Questions Posed in the Call for Papers

On the basis of the above discussion, our opinions on some of the questions posed in the call for papers are as follows.

3.1 What is World Knowledge and What Is Knowledge of Language

Lexical knowledge of a language is relevant for mapping phrases in syntactic structures onto argument positions of a predicate. World knowledge is relevant for representing the distinctions among the senses of particular predicates and arguments which cannot be distinguished by distributional characteristics.

3.2 Cross-Linguistic Evidence for the Specificity of Lexical Semantic Representation

Our experience in medium-scale KBMT projects shows that verb classes defined by linking rules in different languages are not identical. This means that the syntax-driven lexical semantics component of work on a computational application will be different for every

language involved. For example, in Arabic there is a transitivity alternation that is similar to the causative/inchoative alternation (*Someone broke the glass / The glass broke*) in English. However, The English and Arabic transitivity alternations apply to different classes of verbs. The Arabic rule, for instance, applies systematically to verbs of psychological change of state (*excite, please, stun, entertain, confuse, etc.*), but the English rule does not apply to these verbs. Since our goal in developing KBMT systems is to support multilinguality, we preferred to create a single, deeper, interlingual meaning representation which will not be subject to the divergences between verb classes in individual languages (Mitamura, 1989).

4 An Example: Two Semantic Structures

In this section we will show the differences in analyzing a single sentence of English according to each of the two methods of semantic analysis. This sentence has been extracted from a journalistic text, in which it was actually a clause in a compound sentence.

Revenues and earnings will begin to post double digit growth in 1992.

¿From the standpoint of syntax-driven lexical semantics, the desiderata of analysis include preserving argument structure and producing a semantic analysis in such a manner that syntactic realization of predicates and their arguments will be attainable through linking rules. We will concentrate solely on verbs.

Following Jackendoff (1976, 1983, 1990), we will analyze *begin* as a change of circumstance. In other words, *revenues and earnings* enter the circumstance of *posting double digit growth*. For verbs in the change-of-circumstance class, the argument that changes circumstance links with the grammatical function SUBJECT. The circumstance can be realized as a nonfinite complement. We will treat *post* as a conflation whose meaning is similar to “show by posting” and which is a member of the same verb class as *show* or *display* in the sense of *to have an observable property*.⁴ The linking rules for this class link the property with the OBJECT grammatical function and link the object that has the property with the SUBJECT grammatical function.

The LCS for the above example shows two predicates — BEGIN and SHOW-BY-POSTING. Notice that we have avoided the habitual thematic role names for their arguments and adjuncts. Linking to grammatical functions is shown by markers in parentheses.

BEGIN

```

thing-that-changes-circumstance: revenues and earnings (SUBJ)
new-circumstance: SHOW-BY-POSTING (XCOMP)
                    thing-that-has-a-property: revenues and earnings (SUBJ)
                    property: double digit growth (OBJ)
                    time: 1992 (Adjunct)

```

¿From the standpoint of ontology-driven lexical semantics, the desiderata of analysis include a) independence of the syntax of representation from the syntax of the source text, b) explicitly representing all implied meanings which are referred to in the source text — either lexically or through deixis or ellipsis or other means, c) keeping the set of atomic

⁴ A conflation is a combination of two LCSs, A and B, into a new LCS, C. The type of conflation in this example is that C has the meaning of “B by A-ing” and follows the linking rules and subcategorization patterns of B. More on the phenomenon of conflation see in Jackendoff (1990), B. Levin and Rapoport (1988), Talmy (1975, 1985).

meaning representation elements as low as possible (this latter goal does not imply a desire to postulate a small fixed set of meaning primitives). An ontology-driven lexical-semantic representation for the above example can look as follows:⁵

```
clause
head:      increase-1
aspect:
  phase:    begin
  duration: prolonged
  iteration: single
time relative: 'after time of speech'
  absolute: 'during 1992'

increase-1
theme: 'revenues and earnings'
rate: 'greater than 10% and less than 100%'

relation-1
type: temporal-before
from: 'time of speech'
to: 1992
```

Note that the verbal meanings are factored out into propositional meanings of the clause head and the aspectual and temporal meanings. Note also that the clause head is *increase* which means that the meanings of *begin* and *post* got incorporated into other representations — the former gave rise to the value *begin* of the aspectual meaning of the clause. The latter was understood as a functional verb whose only purpose was to allow the English realization of the concept of increasing in a nominalized fashion. Thus, *post* is treated as a collocant of *growth* (see e.g., Mel'čuk, 1981, Nirenburg et al., 1988 or Smadja and McKeown, 1989 for various treatments of collocations). Properties used to describe events and objects (such as *increase*) are predefined in an ontological domain model. Phrases in double quotes are inserted as placeholders, since for the purposes of this comparison their further analysis is immaterial.

5 Conclusion

Our approach to lexical semantics combines the benefits of both syntax-driven and ontology-driven semantics. We have been able to use both types of knowledge in several integrated applications. We plan to use this experience to formulate a distinct lexical-semantic theory based on this work. As a preparatory step toward formulating such a theory we have suggested the structure of ontological domain models, lexicons for particular languages and a text meaning representation language (see Carlson and Nirenburg, 1990; Meyer et al., 1990, Nirenburg and Defrise, 1991a, b).

References

- [1] Brachman, R. 1983. What IS-A and Isn't: An Analysis of Taxonomic Links in Semantic Networks. *Computer*, 16(10), October, 30-36.

⁵In this representation we didn't try to follow any particular knowledge representation formalism.

- [2] Bresnan, J. and J. Kanerva. 1989. Locative Inversion in Chichewa: A Case Study of Factorization in Grammar. *Linguistic Inquiry*, 1-50.
- [3] Carlson, L. and S. Nirenburg. 1990. World Modeling for NLP. CMU CMT Technical Report 90-121.
- [4] Dahlgren, K. and J. McDowell. 1989. Knowledge Representation for Commonsense Reasoning with Text. *Computational Linguistics*, 15, 149-170.
- [5] Fodor, J. A. 1975. **The Language of Thought**. New York: Crowell.
- [6] Grimshaw, J. 1990. **Argument Structure**. Cambridge, MA: MIT Press.
- [7] Hirst, G. 1987. **Semantic Interpretation and Resolution of Ambiguity**. Cambridge University Press.
- [8] Hobbs, J. and R. Moore (eds.) 1985. **Formal Theories of the Commonsense World**. Norwood, NJ: Ablex.
- [9] Horty, J. F., Thomason, R. H., and Touretzky, D. S. (1990) A skeptical theory of inheritance in nonmonotonic semantic nets. *Artificial Intelligence* 42(2-3), 311-348.
- [10] Jackendoff, R. 1976. Toward an Explanatory Semantic Representation. *Linguistic Inquiry*, 89-150.
- [11] Jackendoff, R. 1983. **Semantics and Cognition**. Cambridge, MA: MIT Press.
- [12] Jackendoff, R. 1990. **Semantic Structures**. Cambridge, MA: MIT Press.
- [13] Lenat, D. and R. Guha. **Building Large Knowledge-Based Systems**. Reading, MA: Addison-Wesley.
- [14] Lehnert, W. and M. Ringle (eds.) 1982. **Strategies for Natural Language Processing**. Hillsdale, NJ: Lawrence Erlbaum.
- [15] Levin, B. 1989. English Verbal Diathesis. **Lexicon Project Working Papers 32**. Center for Cognitive Science, MIT.
- [16] Levin, B. and M. Rappaport Hovav. 1990. Wiping the Slate Clean: A Lexical Semantic Exploration. ms. Northwestern University and Bar Ilan University.
- [17] Levin, B. and T. Rapoport. 1988. Lexical Subordination. Proceedings of the 24th Annual Meeting of the Chicago Linguistic Society.
- [18] Levin, L. 1986. **Operations on Lexical Forms: Unaccusative Rules in Germanic Languages**. Ph.D. Dissertation. MIT.
- [19] Levin, L. 1987. Toward a Linking Theory of Relation Changing Rules in LFG. **CSLI Report No. CSLI-87-115, November 1987**, Center for the Study of Language and Information, Stanford, California.
- [20] Mahmoud, A.T. 1989. A Comparative Study of Middle and Inchoative Alternations in Arabic and English. Ph.D. Dissertation. University of Pittsburgh.

- [21] Mel'čuk, I.A. 1981. Meaning-Text Models: A Recent Trend in SOviet Linguistics. *The Annual Review of Anthropology*.
- [22] Meyer, I., B. Onyshkevych and L. Carlson. 1990. Lexicographic Principles and Design for Knowledge-Based Machine Translation. CMU CMT Technical Report 90-118.
- [23] Mitamura, T. 1989. **The Hierarchical Organization of Predicate Frames for Interpretive Mapping in Natural Language Processing**. Ph.D. Dissertation. University of Pittsburgh.
- [24] Nirenburg, S., E. Nyberg, R. McCardell, S. Huffmann, E. Kenschaft, and I. Nirenburg. 1988. DIOGENES-88. Technical Report CMU-CMT-88-107. Center for Machine Translation, Carnegie Mellon University, Pittsburgh, PA.
- [25] Nirenburg, S. and C. Defrise. 1991a. Application-Oriented Computational Semantics. In: R. Johnson and M. Rosner (eds.) **Formal Semantics and Computational Linguistics**. Cambridge University Press (*in press*).
- [26] Nirenburg, S. and C. Defrise. 1991b. Aspects of Text Meaning. In: J. Pustejovsky (ed.). *Semantics and the Lexicon*. Dordrecht, Holland: Kluwer.
- [27] Onyshkevych, B. and S. Nirenburg. 1991. *Lexicon, Ontology and Text Meaning*. This volume.
- [28] Schank, R. 1973. Identification of Conceptualizations Underlying Natural Language. In: R. Schank and K. Colby (eds.), **Computer Models of Language and Thought**. San Francisco: Freeman.
- [29] Smadja, F. and K. McKeown. 1990. Automatically Extracting and Representing Collocations for Language Generation. Proceedings of the 28th Annual Meeting of the ACL.
- [30] Sowa, J. 1984. **Conceptual Structures: Information Processing in Mind and Machine**. Reading, MA: Addison-Wesley.
- [31] Talmy, L. 1975. Semantics and Syntax of Motion. In: J.P. Kimball, ed. **Syntax and Semantics**, 4. New York: Academic Press. 181-238.
- [32] Talmy, L. 1985. Lexicalization Patterns: Semantic Structure in Lexical Forms. In: T. Shopen (ed.) **Language Typology and Syntactic Description**, vol. 3. Cambridge University Press. 57-149.
- [33] Wilks, Y. 1975. A Preferential, Pattern-Seeking Semantics for Natural Language Inference. *Artificial Intelligence*, 6, 144-147.
- [34] Woods, W. 1975. What's in a Link: Foundations of Semantic Networks. In: D. Bobrow and A. Collins (eds.) **Representation and Understanding: Studies in Cognitive Science**. New York: Academic Press.