

The Language of Word Meaning

Pierrette Bouillon and Federica Busa (editors)

(University of Geneva and LingoMotors Inc.)

Cambridge University Press (Studies in natural language processing, edited by Branimir Boguraev), 2001, xvi+387 pp; hardbound, ISBN 0-521-78048-9, \$69.95

Reviewed by

Mari Broman Olsen

Microsoft Corporation

1. Introduction

The papers collected in *The Language of Word Meaning* resemble nothing so much as a holiday celebration in a large, heterogeneous family, with echoes of old feuds, marginally relevant contributions from distant relatives, and fresh stories from recent friends. Although the feuds are entertaining (opening the festivities as Section I, “Linguistic Creativity and the Lexicon”), and the outsider perspectives insightful, from a computational-linguistic perspective many of the most valuable contributions come from guests who have traveled far, and with other companions, before finding common ground at this gathering.

The volume’s title conceals the specificity of its subject: a toast and roast of Generative Lexicon theory (henceforth GL), originally proposed by James Pustejovsky (henceforth JP) (Pustejovsky 1993, 1995, inter alia), who contributed two chapters and the preface. The editors also contribute a joint article, and Busa is coauthor in a second. Others, certainly, would be included in a general discussion on “the language of word meaning.” Indeed (with rare and sometimes confusing exceptions) the discussion assumes GL as the “semantic vocabulary” (p. xv) and focuses on the converse—“the word meaning of language”¹—specifically, how and whether GL’s finite number of generative devices and rules can be used to construct semantic expressions compositionally. Several contributors (Fodor and Lepore, papers in Section II; Vossen, p. 373) raise interesting ancillary questions: how (and whether) to divide word meaning from world knowledge.

Pustejovsky (Chapter 7) and Busa and Bouillon provide especially clear reviews of basic GL concepts. These (and other) surveys would be more useful, however, as a condensed introduction for newcomers (which this review will present with great reduction as part of Section 2). The book claims (according to JP’s preface) that the GL approach to language synthesizes traditions and ideas from ordinary language philosophy (a focus on words and word use), analytic semantics (formalization of rules and types), and generative linguistics (infinite generation of meanings from finite, recursive resources). It claims that GL can tackle empirically difficult problems, both for computational applications and for other formal, even descriptive, systems: specifically, how words vary in context, how new senses emerge, and what the underpinnings are of the systematic mapping of semantic types onto syntactic forms across languages.

¹ Here I’m agreeing with my six-year-old that this would “make more sense” as the title.

The editors begin the volume and each of four sections with short introductions (somewhat confusingly numbered as chapters). The introductions bring harmony and clarity to the four sections (except in Section III, the introduction to which takes issue with Kilgarriff's contribution). Nevertheless, the sections do not clearly form a whole, beyond their family relation to GL, computational lexicography, and/or lexical semantics. JP points out that not all contributors even agree with the generativity of lexical senses and composition. Some who do, but in a different framework (Moravcsik) provide explicit mappings from their representation to GL; others (Asher and Lascarides, Hobbs) leave relationships implicit. Several papers focus on practical natural language processing (NLP) (Section IV and others) and a few on questions of philosophical reality, which, though an interesting diversion, leads to further diffusion of focus.

One frequently wants more explication of criteria. When are new subrelations or attribute values required (cf. Ruimy, Gola, and Monachini)? Why is *book* a physical object in *I am waiting for her next book* (p. 137)? Why does *waiting for the car* anticipate a construction event, but not *repairing the car, washing the car, or waiting for the car to be returned [by a borrower]* (p. 159)? How do we know when senses are different and require different generative mechanisms (p. 367)? Why are some predicted senses rare (pp. 153–154)?

In some chapters, editorial lapses and typos are frequent enough to be distracting: non-English *finish to read* for *finish reading*, and *on one side* for *on one hand*; Schank and Abelson cited as 1997 instead of 1977; unexpanded acronyms (*LCS* by Saint-Dizier, *AFT* by Moravcsik) and multiply expanded acronyms (*JP* by Fodor and Lepore); and a missing abstract (Moravcsik). A few sentences are uninterpretable (the sentence in Yorick Wilks's abstract beginning "It is argued . . ."; a couple of sentences by Piek Vossen (p. 36) on "disjunctive" labels). A shared bibliography would have improved the book's value as a GL reference and reduced the number of pages (all but one paper cites Pustejovsky 1995, for example).

2. Generative Lexicon Theory

GL employs four levels of linguistic representation. ARGUMENT STRUCTURE encodes obligatory and optional arity of predicates. EVENT STRUCTURE describes the aspectual nature of the events, whether states, processes, or transitions (cf. Mourelatos [1978], *inter alia*), and their subevents, if any, as well as the temporal relations between subevents. QUALIA STRUCTURE links the first two representation levels by assigning arguments to participant slots in (sub)events. The four qualia roles, inspired by Moravcsik's interpretations of Aristotle (p. 56), minimally encode the linguistic behavior of lexical items, primarily nominals and predicates with which they compose. The CONSTITUTIVE role of an object represents the relation between it and its parts (and what it is part of, according to Busa and Bouillon), the FORMAL role categorizes it within a larger domain, the TELIC role represents its purpose, and the AGENTIVE role its origins. Finally, LEXICAL INHERITANCE STRUCTURE relates lexical structures to other structures in the type lattice. Generative devices, including type coercion, subselection, and co-composition, connect all four levels in restricted ways, providing for compositional interpretation of words in context (p. 56).

Qualia structure work has proved especially fertile for lexical semantic research, including many clear, data-rich papers in this volume. The qualia roles challenge researchers to consider how far a minimal lexical representation can go toward accounting for semantic productivity, much of which had been assigned to world knowledge. For example, JP suggests (Pustejovsky 1993) that both possible meanings for *I began the book* ('began to read/write') derive from qualia structure, one from the AGENTIVE

(books come into being by writing), and one from the TELIC (books are intended to be read).

Julius M. Moravcsik aligns his fourfold structure (m-factor, s-factor, a-factor, and f-factor) with the four qualia roles above, respectively, to partially decompose metaphors. Salvador Climent argues that the CONSTITUTIVE role accounts for the behavior of Spanish partitives; Patrick Saint-Dizier augments the TELIC role with a small set of rules relating predicates and arguments for more explanatory representation of adjectival modification, as well as some sense variations, metaphors, and metonymies; Jacques Jayez finds qualia structure inadequate to account for the distributional properties of several French synonym classes—for example, to encode the idea that *suggérer* ‘suggest’ permits both animate and inanimate subjects, but the former only if the complement is a “choice” among items in the propositional context. Adam Kilgarriff concludes from the SENSEVAL data (Kilgarriff and Palmer 2000) that qualia structure fails to account for many senses (perhaps metaphorical) not encoded in dictionaries, for which GL would seem to be naturally suited.

3. Section I: Family Feuds

The first section, “Linguistic Creativity and the Lexicon,” considers whether words have internal syntax and inferential meaning (as in GL) or whether they are concepts as particulars with denotational semantics. James McGilvray, in “Chomsky on the Creative Aspect of Language Use and Its Implications for Lexical Semantic Studies,” situates GL as Relation R in Chomskyan syntax, as an internalist semantics, “a branch of syntax, broadly speaking” (p. 19). His leap from the usefulness to the innateness of qualia structure (p. 17) is jarring, as is the sudden appearance of minimalism 14 pages into the 23-page chapter (p. 18). But he clearly and cogently defends Chomsky as grounding a computational approach: obvious to some (p. xii), yet hotly debated in some computational circles.

As they have in the past,² Jerry A. Fodor and Ernie Lepore, in “The Emptiness of the Lexicon: Critical Reflections on J. Pustejovsky’s ‘The Generative Lexicon’”, challenge JP on the possibility of the GL enterprise (or indeed any nonatomic representation structure). Fodor and Lepore (henceforth F&L) lack faith that lexical generalizations are ultimately interesting or powerful and not merely one of “all sorts of ways” in which “the meanings of words can partially overlap” (p. 48). Claiming that the null hypothesis (p. 29) is as “everybody always thought: the lexicons of natural languages are just lists” (p. 44),³ they question whether GL reaches its aims: to account for semantic well-formedness and generativity and to provide evidence that meaning is inferential (rather than atomistic and denotative). JP responds (“Generativity and Explanation in Semantics: A Reply to Fodor and Lepore”) that F&L misread him in substance and detail and are “negative and unconstructive” (p. 51), artificially raising the bar for semantic theory while avoiding the theoretical issues raised by the data. According to F&L, GL is descriptively inadequate, both over- and undergenerating interpretations. JP claims that “the modes of inheritance for concepts associated with linguistic expressions are not arbitrary” (p. 61). F&L say those given are incomplete. F&L’s alternative (*want a beer* → *want to have a beer* → *want to drink a beer* derived by

² See, for example, discussion between Fodor and Weinberg in CUNY (1998), as well as references in the papers by Fodor and Lepore and Wilks in the volume under review.

³ Although, as F&L point out, much effort has been expended studying semantic relations among lexical items, for example, synonymy and antonymy, which are phenomena that they foresee will turn out also to be conceptual rather than lexical.

logical-form rules) should be subjected to the same adequacy standard to which they subject the GL discussion of *bake* (on which JP says, in a footnote on p. 64, that they miss the point anyway).

Wilks stands with JP in “The ‘Fodor’-FODOR Fallacy Bites Back,”⁴ rejecting the F&L position as untenable for analytic reasons (there is no one-to-one mapping between words and meanings) and unhelpful in organizing a useful dictionary, for NLP or otherwise. Although F&L don’t claim “real” representations should be useful, Wilks also disagrees with F&L and JP that an adequate theory of the lexicon should be able to judge semantic well-formedness.

4. Section II: Finding Common Ground

The papers in the second section, “The Syntax of Word Meaning,” refine and gently challenge the vocabulary, framework, and crosslinguistic data analysis required for an explanatorily adequate, language-independent lexical semantics. JP, in “Type Construction and the Logic of Concepts,” a revision of a 1998 paper, discusses well-formedness directly and convincingly, arguing that constraints on concepts (i.e., thought) are revealed through lexicalization strategies. He proposes a tripartite-concept subtyping—natural, functional, and complex, discussed further in Section IV by Busa, Calzolari, and Lenci—and addresses criteria for choosing among competing representations and deciding about feature admission, appropriateness, and composition.

In “Underspecification, Context Selection, and Generativity,” Jayez shows that context provides cues to polysemous word meaning as well as imposing restrictions on synonym selection (French *suggérer* and *attendre* ‘wait’ classes). It is not clear that the data presented are inherently incompatible with the GL framework, as he argues. He also makes the general point that GL practitioners must demonstrate generative adequacy not only against sense enumeration lexicons (the current focus in the literature) but also against other decompositional and underspecified lexical theories.

The next two papers propose extensions to qualia structure. For Bouillon and Busa, in “Qualia and the Structuring of Verb Meaning,” qualia roles are dynamic, varying by speaker and context; for example, the constitutive relation has subtypes (IS PART OF, HAS AS PART, IS IN, LIVES IN, IS A MEMBER OF, HAS AS MEMBER). They argue that the restrictions on the variety of noun phrases that can occur as objects of *attendre* can be derived compositionally: The verb selects an event, and some nominals encode events, for example, creation (such as *journal*, *book*, *car*), possession (all nominals with telic roles, inter alia), and beginning or culminating. Saint-Dizier (“Sense Variation and Lexical Semantics: Generative Operations”) augments the telic role, the “most productive role to novel sense variations” (p. 168).

In “Individuation by Partitive Constructions in Spanish,” Climent argues that the semantics of portions is grounded in the entailment that entities (individuated, mass, etc.) are composed of something. He argues that qualia are syntactic, with compositional operations creating selectionally constrained complex nominals. Through interesting (though constructed) data, in “Event Coreference in Causal Discourses,” Laurence Danlos argues persuasively that causatives can be interpreted from a linguistic perspective with internally complex events, although she rejects the extended event structure for unaccusatives.

⁴ Wilks says F&L claim to be part of informational role semantics (IRS), as he is. This is confusing, since in their chapter they express doubt that IRS can be sustained, challenge the cogency of arguments that take it as a premise, and go on to address JP, who purports to provide an argument for the complexity of lexical entries that do not presuppose IRS (p. 29).

5. Section III: Forbears and New Friends

The gratingly titled third section, “Interfacing the Lexicon” (with what?), examines metonymy and metaphor in the GL framework. The stimulating studies in this section decompose structures generally considered lexically atomic. In “Metaphor, Creative Understanding, and the Generative Lexicon,” Moravcsik uses GL’s polysemy account (whose qualia spring, in part, from Moravcsik’s work) to uncover the lexical information structure by means of metaphor. Moravcsik argues that simile also can be treated “as a literal statement” (p. 260) of comparison.

Nicholas Asher and Alex Lascarides examine “Metaphor in Discourse” in terms of lexical semantic structures, inspired by GL and others (Copestake and Briscoe 1995), focusing on change-of-location verbs, metaphorical extension of properties from objects to persons, and the relation between metaphorical interpretation and discourse structure. They constrain interpretations by means of lexical rules that restrict meaning shifts and a discourse structure theory in which truth conditions can be evaluated.

Jerry Hobbs presents his interpretation-as-abduction framework in “Syntax and Metonymy.” He accounts for Nunberg’s (1995) metonymic data (*I’m parked out back*) by means of a set of coercion relations, selected on the basis of saliency by least-cost abductive interpretation. He extends his account to a wide range of intriguing data, old and new, including extraposed modifiers (*A jolly old man arrived with an armload of presents*), Bolinger’s (1988) ataxis (*John smokes an occasional cigarette = John occasionally smokes cigarettes*), container nouns (cf. Jayez), and disguised small clauses (*I have a sore throat = My throat is sore*).

Kilgarriff, in “Generative Lexicon Meets Corpus Data: The Case of Nonstandard Word Uses,” brings to the discussion a refreshing quantitative, corpus linguistics perspective. He asks to what extent GL is able to generate, for a sample set of words, senses found in corpora but not represented in a dictionary (“nonstandard” senses). As mentioned above, he argues that contextual sense variation is too large and unpredictable for such a model of lexical structures. Some might question his strategy for determining nonstandard senses and his rigorous interpretation of GL mechanisms, but his methodology is clear, reproducible, and a fair challenge to GL, potentially accounted for by Bouillon and Busa’s expansion of the qualia relations, for example.

6. Section IV: Practical Matters

The fourth section, “Building Resources,” chronicles practical experiences for a range of NLP applications within the SIMPLE and EuroWordNet frameworks. The papers in this section, as the introduction states, leave open the question of whether they are useful or usable; they furthermore neglect to evaluate whether they are practical to construct: Most involve significant manual work.

In “Generative Lexicon and the SIMPLE Model: Developing Semantic Resources for NLP,” Busa, Nicoletta Calzolari, and Alessandro Lenci highlight the need to capture richness of language (apparently independent of applications) within a testable model for building large-scale resources. They employ qualia roles for structuring the semantic/conceptual types, assuming that words vary in complexity (expanding on Pustejovsky in Section II). Nilda Ruimy, Elisabetta Gola, and Monica Monachini, in “Lexicography Informs Lexical Semantics: The SIMPLE Experience,” describe how they use GL to categorize concepts, without defending or critiquing the theory or their particular attribute-value assignments. They argue that encoding their methodology serves to explicate the theory.

Piek Vossen describes his work on “Condensed Meaning in EuroWordNet,” focusing on the challenge of developing GL types of relations from resources that are sense-enumerative. He sets forth how EuroWordNet links semantic representations across languages via an Inter-Lingual-Index that connects their structures.

7. Conclusion

The Language of Word Meaning offers glimpses into the variety of work spawned by GL: philosophical, computational, theoretical, and multilingual. Section introductions and GL overviews by Pustejovsky (Chapter 7) and Busa and Bouillon lay out important themes in computational and theoretical lexical semantics. Other chapters offer practical, often entertaining exemplars of lexical work unfolding—like family celebrations—somehow, noisily, and with great diversity.

References

- Bolinger, Dwight. 1988. Ataxis. In Rokko Linguistic Society, editors, *Gendai no Gengo Kenkyu* [Linguistics Today], Kinseido, Tokyo, 1–17.
- Copestake, Ann and Ted Briscoe. 1995. Semi-productive polysemy and sense extension. *Journal of Semantics*, 12(1):15–67. Reprinted in James Pustejovsky and Branimir Boguraev, editors, *Lexical Semantics: The Problem of Polysemy*, Oxford University Press, 1996, 15–67.
- CUNY. 1998. *Proceedings of the 11th CUNY Conference on Human Sentence Processing*, Rutgers University, New Brunswick, NJ.
- Fodor, Jerrold. 1970. Three reasons for not deriving *kill* from *cause to die*. *Linguistic Inquiry*, 1:429–438.
- Kilgarriff, Adam and Martha Palmer, editors. 2000. Special Issue on SENSEVAL: Evaluating Word Sense Disambiguation Programs. *Computers and the Humanities*, 34.
- Mourelatos, Alexander. 1978. Events, processes, and states. *Linguistics and Philosophy*, 2:415–434. Also published in Philip Tedeschi and Annie Zaenen, editors, *Syntax and Semantics 14: Tense and Aspect*, Academic Press, 1981, 191–212.
- Nunberg, Geoffrey. 1995. Transfers of meaning. *Journal of Semantics*, 12:109–132.
- Pustejovsky, James. 1993. Type coercion and lexical selection. In James Pustejovsky, editor, *Semantics and the Lexicon*. Kluwer Academic, pages 73–96.
- Pustejovsky, James. 1995. *The Generative Lexicon*. MIT Press.

Mari Broman Olsen has worked as lexical semanticist in Microsoft Corporation’s Natural Language Group since 1999. From 1996 to 1999, she was a postdoctoral fellow at the University of Maryland Institute for Advanced Computer Studies, researching interlingual lexical representation for Chinese–English machine translation, inter alia. Her Ph.D. is from Northwestern University (1994); her dissertation, *A Semantic and Pragmatic Model of Lexical and Grammatical Aspect*, was published by Garland Press. Olsen’s address is: Microsoft Corporation, One Microsoft Way, Redmond, WA 98052; e-mail: molsen@microsoft.com.