

Book Reviews

Tree Adjoining Grammars: Formalisms, Linguistic Analysis and Processing

Anne Abeillé and Owen Rambow (editors)
(Université de Paris VII and AT&T Labs–Research)

Stanford, CA: CSLI Publications
(distributed by the University of
Chicago Press), 2000, vii+478 pp;
hardbound, ISBN 1-57586-251-4, \$64.95;
paperbound, ISBN 1-57586-252-2, \$24.95

Reviewed by
Geoffrey K. Pullum
University of California, Santa Cruz

Take a finite set \mathcal{T} of trees, and close it under the operation of *substitution*—replacing daughterless nonterminals by other trees in \mathcal{T} whose root node label matches the nonterminal. The set of all terminal strings of trees in the resultant tree-set will be a context-free language (CFL), and for every CFL there will be such a finite set of trees that generates it under substitution. For an elegant formalization of CFLs in these terms, see Rogers (1999, pages 25–26).

Now assume an additional operation: besides substituting trees for daughterless nonterminals, you can also squeeze new material into the middle of a tree, substituting it for a nonterminal node that has daughters. More precisely, assume a finite set \mathcal{A} of insertable trees that share a special form: each $A \in \mathcal{A}$ has a single node on its frontier, known as A 's foot node, that has the same label as A 's root. Squeezing some $A \in \mathcal{A}$ into a tree $T \in \mathcal{T}$ means replacing a node n in T so that n 's mother becomes the mother of A 's root, and everything dominated by n comes to be dominated by A 's foot node.

Closing \mathcal{T} under the operation of squeezing in new material from \mathcal{A} in this manner (an operation called tree adjunction) yields a tree-set of which the set of all terminal strings will be a tree adjoining language (TAL), and for every TAL there will be an appropriate pair $\langle \mathcal{T}, \mathcal{A} \rangle$. Such a pair (called a tree adjunct grammar in Joshi, Levy, and Takahashi [1975]) is known today as a tree adjoining grammar (TAG).

The research program on TAGs that Aravind Joshi has led since 1975 is perhaps the most interesting and significant research program in formal language theory of the last 40 years. General linguists have clearly underrated it, though computational linguists have in general kept more closely in touch with it. The TALs are a mathematically natural class with closure and decision properties very similar to those of the CFLs, including a polynomial-time recognition problem. Several independent but equivalent characterizations of the class have been discovered: Vijay-shanker and Weir (1994) present a weak equivalence result for head grammars (Pollard 1984; Roach 1987), linear indexed grammars (Duske and Parchmann 1984; Gazdar 1988), and combinatorial categorial grammars (Steedman 1986). Additionally, Vijayashanker (1988) gave a characterization in terms of embedded pushdown automata; and Rogers (1998) gives a new model-theoretic characterization in terms of linearized terminal strings of three-dimensional tree models of monadic second-order logic formulae.

Moreover, the descriptive capabilities of TAGs make them as plausible a theory of syntax for natural languages as has ever emerged from formal language theory. The ways in which TAGs exceed the descriptive capacity of context-free grammars seems remarkably close to what one would want in a theory of syntax sculpted for human languages.

However, like relational grammar in the 1970s and optimality theory in the 1990s, the theory of TAGs has never had the book-length exposition it deserves: there has been no coherent and comprehensive published monograph by the original developers that gives an integrated account of the framework and convincing examples of its application to a well-known language.

Tree Adjoining Grammars, edited by Anne Abeillé and Owen Rambow, does not entirely fill that gap, though it goes some of the way. It is a refereed collection of papers that were presented in earlier versions at an international workshop on TAGs and related formalisms in Paris in 1994. The editors contribute a substantial introduction, and the 18 other chapters are grouped into three sections: (1) "Formalisms," (2) "Linguistic Analysis," and (3) "Processing."

I cannot here summarize or critique all the papers in this rich collection, but I will comment briefly on a few of the papers that I think would on their own justify the purchase of the book (especially at CSLI's attractively low price of under \$25 for the paperback).

Chapter 1 is an impressive 68-page essay by Abeillé and Rambow. An expansion of this essay with fuller exemplification and more precision in formulation, perhaps coauthored with Joshi, might have made an independent monograph on TAGs that would have filled the gap referred to above. But instead this lengthy essay is just an extended introduction to an anthology. It is not without flaws. For example, its introduction of the crucial notion "lexicalized" is remarkably casual and does not really permit one to discern what the definition is. (My understanding is that in a lexicalized TAG, each of the trees in \mathcal{T} must contain one and only one terminal symbol [word in the dictionary]. But in that case, the diagrams in this chapter never really give an example of a lexicalized TAG, which is odd, since the authors take this concept to be centrally important.) That said, however, the amount of work this chapter represents is substantial, and the standard of exposition is mostly high. It does a lot more than the usual summary-of-the-rest-of-the-book that is typical for an editorial introduction.

Chapter 2, by Roger Evans, Gerald Gazdar, and David Weir, has an oversubtle title that is ruined by capitalization of significant words on the chapter title page and in the running heads. The title should read as follows:

"Lexical Rules" are just lexical rules.

The thesis is that the theoretical machinery needed to state the sort of lexical rules that everybody assumes (like the one that assigns irregular plurals in *-i* to certain Latinate nouns ending in *-us*), properly understood, suffices to take over all the work of the special class of generative "Lexical Rules" that many frameworks recognize—the rules for covering phenomena like the active/passive relation between verb subcategorization frames. A full consideration of what is needed to state the former type of lexical generalization in the representation language DATR reveals that such mechanisms can also handle the (apparently) heavier stuff. Building on earlier work by the authors (Evans, Gazdar, and Weir 1995), this elegant paper shows how to state lexicalized TAGs compactly and nonredundantly using DATR, and how to express a variety of so-called Lexical Rules without positing any additional devices.

Chapter 6, “Complexity of Scrambling: A New Twist to the Competence-Performance Distinction” by Aravind K. Joshi, Tilman Becker, and Owen Rambow, deserves to become known as a classic. It proposes that certain limitations on scrambling (re-ordering of clausal constituents) in German should not be treated as performance limits (in the way the limits on center-embedding in English usually are) because there is a way to make them follow from a grammar formalism (namely, TAGs), and this provides a better explanation than positing a performance restriction stemming from some unknown psychological or neurological basis. In other words, the paper advocates letting the syntactic theory decide: whether a theory is available that will draw a certain syntactic distinction should be a relevant factor in deciding when and whether performance limitations are to be invoked. The argument is clever, convincing, and quite surprising.

Chapter 12, “Implications of Binding for Lexicalized Grammars” by Mark Steedman, is a characteristically wide-ranging and interesting look at what binding phenomena mean for transformational grammar, generalized phrase structure grammar, head-driven phrase structure grammar, TAGs, and various types of categorial grammar (recall that Steedman’s combinatory categorial grammar is weakly equivalent to TAGs).

Many other papers in the volume will repay study. Among the less technical are Robert Frank’s paper speculating on children’s progress toward increasing syntactic complexity during language acquisition (Chapter 3), which is moderately interesting but does not go much beyond the suggestion that comparing acquisition time and processing load for various constructions might be a good idea, and the paper by Beth Ann Hockey and Heather Mateyak on the semantic features that influence the sequencing of determiners in English (Chapter 9), which does not make much essential use of TAGs. But others are highly (even indigestibly) technical studies of various theoretical modifications of or alternatives to TAGs (e.g., Gisela Pitsch’s comparison of TAGs with hyperedge replacement grammars, Chapter 7). Most of the examinations of parsing and implementation issues (Chapters 13–19) are fairly demanding.

This is a valuable book, and I am glad to have it. But it is my duty as reviewer to express a small grumble about it. This book is not a credit to the editorial profession. The bibliographies to the chapters (some 30 pages altogether) are not harmonized in style (e.g., with respect to capitalization) and are not collated at the end of the book (which makes for some wasteful duplication). The index is simply unacceptable: it will not enable scholars to find in this book the things they are looking for. And the text contains many misprints and formatting errors. I noted: “probelm” for “problem” (page 51); “explicitely” for “explicitly” (page 147); “theXP” for “the XP” (page 169); “connectivesor” for “connectives or” and “markerscan” for “markers can” (page 249); “alors.” for “alors.” (page 253); “bindingpossibilities” for “binding possibilities.” (page 283); “analysisas such” for “analysis as such” and several other such errors (page 284, bottom); “gapsare” for “gaps are” and several other such errors (page 298, bottom); “slowness” for “slowness” (page 324); “asemantic graph ? an answer” for “a semantic graph? An answer” (page 324); “Gazdar, G. G.” for “Gazdar, G.” (page 471); and so on (this list is not exhaustive).

There are typos in most books, of course; accurate proofreading is an arduous job. But this book falls below what might reasonably be expected. To see an extraneous paragraph break caused when \LaTeX hit a double line break in the source file (see page 77) suggests that some parts of the book were hardly even looked at in final form, let alone proofread with care, by either the editors or the publisher’s staff.

If books are going to be produced via ready-to-run \LaTeX files, and volume editors do not take their jobs seriously, it bodes ill for the future of books. \LaTeX makes beautiful

pages (Donald Knuth and Leslie Lamport did their jobs), but it can't spell, and it can't insert or delete word breaks or line breaks in the source file. The jobs of copy editors and compositors and proofreaders still have to be done (the production of this journal by The MIT Press still involves a copy editor, a proofreader, and an expert L^AT_EX wrangler over and above the editors). I am sorry to say that the editors of this generally interesting and useful book have let their readers down.

References

- Duske, Jürgen and Rainer Parchmann. 1984. Linear indexed languages. *Theoretical Computer Science*, 32:47–60.
- Evans, Roger, Gerald Gazdar, and David Weir. 1995. Encoding lexicalized tree adjoining grammars with a nonmonotonic inheritance hierarchy. In *Proceedings of the 33rd Annual Meeting of the Association for Computational Linguistics*, pages 77–84.
- Gazdar, Gerald. 1988. Applicability of indexed grammars to natural languages. In Uwe Reyle and Christian Rohrer, editors, *Natural Language Parsing and Linguistic Theories*. D. Reidel, Dordrecht, pages 69–94.
- Joshi, Aravind K., Leon S. Levy, and Masako Takahashi. 1975. Tree adjunct grammars. *Journal of Computing and System Sciences*, 19:136–163.
- Pollard, Carl. 1984. *Generalized Context-Free Grammars, Head Grammars and Natural Language*. Ph.D. thesis, Stanford University.
- Roach, Kelly. 1987. Formal properties of head grammars. In Alexis Manaster-Ramer, editor, *Mathematics of Language*. John Benjamins, Amsterdam, pages 293–348.
- Rogers, James. 1998. A descriptive characterization of tree-adjoining languages. In *Proceedings of the 17th International Conference on Computational Linguistics (COLING'98) and the 36th Annual Meeting of the Association for Computational Linguistics (ACL'98)*, pages 117–121.
- Rogers, James. 1999. The descriptive complexity of generalized local sets. In Hans-Peter Kolb and Uwe Mönnich, editors, *The Mathematics of Syntactic Structure: Trees and Their Logics*. (Studies in Generative Grammar, 44.) Mouton de Gruyter, Berlin, pages 21–40.
- Steedman, Mark. 1986. Combinators and grammars. In Richard Oehrle, Emmon Bach, and Deirdre Wheeler, editors, *Categorial Grammars and Natural Language Structures*. Foris, Dordrecht, pages 417–442.
- Vijayashanker, K. 1988. *A Study of Tree Adjoining Grammars*. Ph.D thesis, University of Pennsylvania.
- Vijay-shanker, K. and David J. Weir. 1994. The equivalence of four extensions of context-free grammars. *Mathematical Systems Theory*, 27:511–546.

Geoffrey K. Pullum is professor of linguistics at the University of California, Santa Cruz, where his teaching ranges from a linguistics graduate course in mathematical foundations of linguistics to a computer science freshman course on the Unix operating system. He is coauthor with Rodney Huddleston of a forthcoming book entitled *The Cambridge Grammar of the English Language* (Cambridge University Press). Pullum's e-mail address is pullum@ling.ucsc.edu; URL: <http://ling.ucsc.edu/~pullum>.

The Theory and Practice of Discourse Parsing and Summarization

Daniel Marcu

(Information Sciences Institute, University of Southern California)

Cambridge, MA: The MIT Press, 2000,
xix+248 pp; hardbound, ISBN
0-262-13372-5, \$39.95

Reviewed by

Udo Hahn

Albert-Ludwigs-Universität Freiburg

Marcu's monograph is based on his Ph.D. thesis—research carried out at the Department of Computer Science, University of Toronto—and subsequent work conducted at the Information Sciences Institute, University of Southern California. It argues for the idea that discourse/rhetorical relations that connect text spans of various length can be computed *without* a complete semantic analysis of sentences that make up these text segments. As an alternative, a formal specification of admissible text structures is provided, which constrains the range of possible semantic and functional connections between text spans and imposes strict well-formedness conditions on valid discourse structures. For effectively computing these text structures, mainly surface-oriented lexical cues and shallow text-parsing techniques are used. Complementary to these formal and computational considerations, Marcu reports on various evaluations, both intrinsic and extrinsic, in order to assess the strengths and weaknesses of his approach and the generality of the principles it is based on. These experiments were mostly carried out on *Scientific American*, TREC, MUC, *Wall Street Journal*, and Brown corpora.

The book consists of three main parts. In the first part, linguistic and formal properties of coherent texts are discussed, with a focus on high-level discourse structures. This theoretical framework serves, in the second part, as the background for developing discourse structure parsing algorithms that compute rhetorical relations in real-world free texts. The benefits of such algorithms for building a high-performance text summarization system are dealt with in the third part.

In the first part, the author factors out a set of assumptions that are common to prominent approaches to discourse structure. So, consensus has been reached that texts can be segmented into nonoverlapping, elementary textual units, that discourse relations of different types link (elementary and complex) textual units of various sizes, that some textual units are more important to the writer's communicative intentions and goals than others, and that trees are a good approximation of the abstract structure of most texts. These considerations lead to a compositionality criterion that requires that discourse relations that link two large text spans can be explained by discourse relations that hold between at least two of the most salient text units of the constituent spans. This notion then forms the basis for a first-order logic axiomatization that captures formal properties of valid text structures. Although this formalization is independent of the set of rhetorical relations actually considered, it yields, by proper relation instantiation, a formal characterization of the structural properties that are specific to Rhetorical Structure Theory (RST) (Mann and Thompson 1988). Building on these formal considerations, the author discusses three (nonincremental) algorithmic paradigms that compute some or all valid discourse structures of a text. Two of them employ model-theoretic techniques and encode the problem of text-structure

derivation as a classical constraint satisfaction problem and as a propositional satisfiability problem. The other one is grammar-based and builds on a proof theory for solving the text-structure derivation problem (demonstrated to be sound and complete with respect to the given logical formalization). The performance of these algorithms is compared empirically on a benchmark of eight manually encoded text-structure derivation problems.

Marcu uses logic to distinguish between discourse structures that are valid and those that are not, so that all valid discourse structures of a text can be determined. In the second part of the monograph, attention then shifts to alternative approaches to deriving valid discourse structures. The first approach relies primarily on discourse markers for shallow rhetorical parsing and employs, as a result of an in-depth corpus analysis, manually designed rules covering more than 450 English cue phrases such as *because*, *however*, and *in addition*, as well as punctuation marks. The second approach adds to plain discourse markers knowledge of surface-oriented lexical co-occurrence data, syntactic criteria (such as part-of-speech categories), and lexical similarity measures based on semantic relation information in order to identify text segments and their rhetorical organization. Given this knowledge-richer setting, discourse parsing rules were automatically derived by applying machine learning techniques (the C4.5 decision-tree algorithm) to data obtained from three corpora of manually annotated discourse trees. All these approaches are meticulously and lucidly described by providing various algorithm schemata for relevant computation steps. Empirical studies are then concerned with the role that discourse markers play in properly segmenting texts into elementary text units and in signaling rhetorical relations that hold between the text segments they connect. The correctness of the discourse trees built by the parser is judged intrinsically, by comparing automatically derived trees with ones that have been built manually, as well as extrinsically, by evaluating the impact automatically derived discourse trees have on properly solving natural language processing problems such as the summarization of texts.

In the third part of the book, the utility of computing discourse structures is empirically assessed in the context of such a text summarization (i.e., extraction) task. The approach advocated by Marcu is readily applicable to this problem, since the representation structures it yields offer implicit content salience orderings in terms of the hierarchical tree structure and the distinction of important information contained in the nucleus and less-important information contained in the satellite portion of text spans, all of which are of immediate relevance for summarization purposes. The main hypothesis to be confirmed is whether or not discourse structures can be successfully exploited in a practical summarization setting. In a methodological experiment, evidence is gathered that text structures such as those mentioned above indeed effectively contribute to identifying the most important units of a text. A discourse-structure-based summarization algorithm that builds on these principles implements a simple salience metric that interprets the tree structure generated by the simple cue-phrase-based text-structure parser. A comparative evaluation reveals that this approach significantly outperforms two baseline algorithms (lead sentence and random sentence selection) and Microsoft's Office 97 summarizer. Considering the structure of discourse to be the paramount factor in determining salience, and incorporating a variety of additional position-, title-, text-tree-, and lexically-based summarization heuristics, a simple GSAT-style learning mechanism is presented that optimizes a linear combination of seven single salience metrics (in terms of combined recall and precision). This way, a significant increase in the performance of the discourse-based summarizer is achieved (yet parameter tuning is clearly dependent on the given text genre and compression rate!).

Marcu's monograph presents a cornerstone in the computational treatment of texts. It has formal merits, as it provides a model-theoretic framework for the study of text coherence structures, in general, and the study of RST, in particular. It has computational merits, as it provides alternative ways of deriving text-structure descriptions automatically and inexpensively (i.e., avoiding full, in-depth text understanding) and distinguishes, given the *a priori* axiomatization, valid text structures from invalid ones. It has methodological merits, as it incorporates machine learning techniques for automatically acquiring the rules needed for discourse parsing and discourse-structure-driven summarization. Finally, it has empirical merits, as algorithms are tested and validated under different experimental conditions.

Marcu also frankly admits that his work ignores the wealth of linguistic constructs that have been shown to be important in text understanding. Such phenomena include focus, topicality, cohesion and reference, pragmatics, and so on. Hence, the notion of validity being proposed is a constrained one, and it has to be weighed carefully against the notion of adequacy and expressiveness of the representation structures derivable therefrom. Still, the author claims that these phenomena can be couched in his formal framework as well. Additionally, one might mention the crucial role of domain-knowledge-dependent inferences and their interaction with building text structures in the absence of explicit cue phrases. Further open issues are the granularity of the text units that span rhetorical relations (e.g., the phrasal as opposed to the clausal or "clause-like" level) and the impact of the text genres under scrutiny. Finally, the dependence on basic assumptions and constructs underlying RST, despite the author's attempt to abstract away from it as much as possible, might be more prevalent than is acknowledged.

The book spans a wide variety of issues in a well-structured, reader-friendly way, and it is easy to understand even in its technical passages. Hence, it can be highly recommended for graduate courses on text analysis. Students are given an outstanding example of the current research paradigm of computational linguistics, which includes formal, algorithmic, methodological, and empirical contributions. And they also may learn how scientific results can be communicated in a rigorous though comprehensible manner.

Reference

- Mann, William C. and Sandra A. Thompson.
1988. Rhetorical Structure Theory: Toward
a functional theory of text organization.
Text, 8(3):243–281.

Udo Hahn is a professor of computational linguistics at Albert-Ludwigs-Universität Freiburg, Germany. His methodological interests include text parsing, knowledge and discourse representation, and learning from texts. He has worked mainly on text analysis applications such as text summarization, knowledge extraction and text mining, and document retrieval. Hahn can be contacted via www.coling.uni-freiburg.de/~hahn.

Natural Language Semantics

Keith Allan

(Monash University)

Oxford: Blackwell Publishers, 2001,
xix+529 pp; hardbound, ISBN
0-631-19296-4, \$73.95, £65.00;
paperbound, ISBN 0-631-19297-2,
\$41.95, £17.99

Reviewed by

Rodger Kibble

Goldsmiths College

This is a large volume, and it contains multitudes. Semantics is construed in a broad sense as the study of how meaning is communicated through the medium of language in a social context, taking account of inferences the hearer is expected to make on the basis of such factors as linguistic knowledge per se, context and “co-text,” encyclopedic knowledge, conventions of politeness and cooperative behavior, and the relative social status of speaker and hearer. The book ranges over a variety of approaches that have addressed these issues, including philosophy of language, lexicography, formal (logic-based) and cognitive semantics, frame-based knowledge representation, pragmatics, and anthropology. However, the result is more than a catalogue of theoretical tools and frameworks; throughout the book, Allan keeps in view an underlying philosophy that “meaning is cognitively and functionally motivated.”

Chapters 1 and 2 introduce fundamental notions such as sense and reference, extension and intension, compositionality, and speech acts. Chapters 3–5 deal with aspects of lexical semantics: Chapter 3 concerns the structure and content of lexical entries, Chapter 4 investigates the extent to which individual morphemes can be assigned semantic interpretations, and Chapter 5 contains an illuminating discussion of aspects of nonliteral word meaning such as connotation, euphemism, dysphemism, and jargon. Chapters 6 and 7 introduce the formal apparatus of propositional and predicate logic and the lambda calculus, and discuss notions of consequence such as semantic entailment and conversational and conventional implicatures. Chapters 8–10 are concerned with “cognitive and functional approaches to semantics,” that is, approaches whose theoretical constructs are claimed to have some form of “psychological reality” or are motivated in terms of their “communicative functions.” These chapters review topics such as frames and scripts, componential analysis, classifiers, color categories across languages, prototypes, and stereotypes. Chapters 11–13 address various issues in clausal and nominal semantics. Chapter 11 concerns modality, tense, and thematic roles, whereas Chapter 12 discusses different approaches to the semantics of verbs and other predicates. Chapter 13 grapples with some of the intricacies of quantification, number, and countability in English noun phrases, using generalized quantifier theory and a variant of ensemble theory. As far as I know, this is the first published tutorial account of the latter, which has previously only been accessible to students in Harry Bunt’s rather challenging monograph (1985). We are told that Chapters 11–13 “demonstrate the application of formal methods of semantic analysis to a corpus of data.” This section is likely to disappoint computational linguists, who will understand the term *corpus* in a different way, since the data in these chapters consist of a series of single-sentence (and mostly single-clause) examples apparently constructed by the author.

The author is clearly in sympathy with the cognitive semantics school, which claims to uncover “psychologically real” structures and processes involved in language use. It’s not always clear to me from Allan’s account what the various claims for psychological reality amount to (such as that for “linguistic categories, semantic fields, frames and the like,” page 288). The methodology displayed tends to follow the standard practice of linguistics textbooks in postulating abstract analyses of examples constructed by the analyst, the reader being invited to share the analyst’s intuitions about their acceptability and interpretation; there is little appeal to experimental or neurological evidence, for example. There are some fascinating discussions of various senses of words such as *back* (pages 289ff.) and *over* (pages 330–331), extended from their basic senses that are presumed to be rooted in direct physical perception. However, these do not give rise to productive procedures that could reliably generate extended senses for equivalent words in other languages, for instance, or other words denoting physical relations or body parts. An attempt to map out some common ground between the cognitive and formal approaches is far from convincing (pages 288–289, emphasis added):

(First premise) Formal representations are created by human minds and are interpretable by human minds. *Therefore*, they have cognitive reality . . .

(Second premise) The informal metalanguages of the cognitivists . . . are creations of deliberate, consciously contrived artifice, just as much as any formal metalanguages are.

(Conclusion) Formalists, cognitivists and functionalists all use contrived metalanguages that have cognitive reality.

As an introduction to *formal* semantics, this book does not supersede established classics such as Gamut (1991). Definitions are sometimes unsatisfactory, effectively substituting one imprecise term for another, as when Grice’s (1975) conventional implicature (CI) is defined as “implies . . . but does not entail” (page 189). The examples of CI that are offered include these: *all gold* implies that the ensemble of gold is nonempty (page 437) and *four eggs* implies *at least two eggs* (page 189). But these are surely different phenomena, the former being a defeasible convention and the latter an arithmetical consequence of the meanings of *four* and *two*. Likewise, in the chapter on quantifiers, *few students* is glossed as “very-much-less than” all contextually relevant students (page 433), though how much less counts as “very-much-less than” itself depends on context and assumptions about prior expectations (Moxey and Sanford 1993).

Every computational linguist should own at least one semantics textbook. Allan’s book stands apart from many other texts in the way it conveys a real sense of the variety and fecundity of language as spoken by living, breathing human beings, rather than as a source of intriguing logicophilosophical puzzles. Nonspecialists will certainly find it an informative, albeit uneven, conspectus of paradigms and areas of inquiry in linguistic semantics, with the bonus that it is actually fun to read.

References

- Bunt, Harry. 1985. *Mass Terms and Model-Theoretic Semantics*. Cambridge University Press, Cambridge, UK.
- Gamut, L. T. F. 1991. *Language, Logic, and Meaning* (2 vols.). University of Chicago Press, Chicago.
- Grice, H. P. 1975. Logic and conversation. In Peter Cole and Jerry L. Morgan, editors, *Syntax and Semantics 3: Speech Acts*. Academic Press, New York, pages 41–58.
- Moxey, Linda and Anthony J. Sanford. 1993. *Communicating Quantities*. Lawrence Erlbaum Associates, Hillsdale, NJ.

Rodger Kibble lectures at Goldsmiths College, University of London, and has research interests in formal and computational approaches to natural language semantics. His address is Department of Mathematical and Computing Sciences, Goldsmiths College, London SE14 6NW, UK; e-mail: R.Kibble@gold.ac.uk.

Intonation: Analysis, Modelling and Technology

Antonis Botinis (editor)

(University of Skövde and University of Athens)

Dordrecht: Kluwer Academic Publishers (Text, speech and language technology, edited by Nancy Ide and Jean Véronis, volume 15), 2000, ix+396 pp; hardbound, ISBN 0-7923-6605-0, \$156.00, £90.00, €145.00; paperbound, ISBN 0-7923-6723-5, \$54.00, £34.00, €50.00

Reviewed by

Martine Grice

Saarland University

The study of intonation is an expanding field, extending beyond core linguistic disciplines such as syntax, semantics, and pragmatics into areas as wide-ranging as psycholinguistics, neurolinguistics, discourse analysis, and emotion research. Intonation is also currently the prime focus of attention in speech synthesis research and is rapidly gaining ground in speech recognition. This expansion has been reflected in a number of workshops and conferences devoted solely to intonation and its interfaces, one of which was the ESCA Workshop on Intonation in Athens in September 1997. The volume reviewed here is one of two collections of papers based on contributions to this workshop. The other collection has appeared as a special issue of *Speech Communication* (33(4), March 2001).

The book consists of four main sections—Prominence and Focus, Boundaries and Discourse, Intonation Modelling, and Intonation Technology—along with an introduction (Antonis Botinis) and a historical overview (Mario Rossi). In this review, I shall concentrate on papers that relate intonation to semantic, pragmatic, or discourse functions and leave the papers dealing solely with speech or phonetics for a review in a journal specializing in those areas.

In the Prominence and Focus section, Julia Hirschberg and Cinzia Avesani's paper, "Prosodic Disambiguation in English and Italian," investigates to what degree speakers of English and Italian use intonational means to disambiguate semantically and syntactically ambiguous sentences. The authors found that, with the exception of quantifier scope, semantic ambiguities were generally more clearly disambiguated than syntactic ones. This was true for both languages. Regarding the semantic ambiguities, the two languages had similar strategies: the scope of negation was disambiguated by phrasing, and the differences in scope of focus-sensitive operators were distinguished by means of pitch accent placement. The authors note that although speakers were aware of the distinctions, they often produced a neutral rendition that would be felicitous for either interpretation, presumably because they read the tokens within a context that already resolved the ambiguity. We learn from this that although context is necessary for eliciting the correct reading, it may at the same time dispense with the reason to disambiguate.

In the Boundaries and Discourse section, all three papers have something to offer for readers of *Computational Linguistics*. Vincent van Heuven and Judith Haan's "Phonetic Correlates of Statement versus Question Intonation in Dutch" is based on both

production and perception experiments. The authors show that although questions as a category have a number of acoustic properties that clearly distinguish them from statements, each question type has a distinct profile of its own in terms of F_0 (fundamental frequency, the perceptual correlate of which is pitch). These question types are *wh*-questions (lexical and syntactic marking), yes-no questions (which in Dutch are marked syntactically), and declarative questions (syntactically indistinguishable from statements). One major cue for the perception of questions is a sentence-final rise in pitch that was never found in statements. The smaller the number of lexicosyntactic indicators as to interrogativity, the higher the rise in pitch and the greater the incidence of such a rise. Other cues include the pitch range, height, and overall shape (e.g., downward or upward trends) across the whole sentence.

Monique van Donzel and Florian Koopmans-van Beinum's "Pitch Movements and Information Structure in Spontaneous Dutch Discourse" confirms previous findings that new information is more often accented than inferrable information. The authors establish the following hierarchy of accentability: new information > inferrable information > verbs > modifiers > discourse markers and evoked information. Discourse boundaries, assigned on the basis of a discourse model developed by the authors in earlier work, are realized with rising pitch (labeled as nonfinal by naive listeners) more often than previously reported. The authors also show that pitch height depends on neither newness nor the type of discourse boundary. Speakers varied a great deal with regard to how often they marked discourse boundaries with pitch variation as opposed to, say, pausing. However, despite these realizational differences, naive listeners perceived prominences and boundaries to a comparable extent across speakers, indicating that they are flexible enough to adapt their perceptual criteria to the current speaker.

Anne Wichman, Jill House, and Toni Rietveld's "Discourse Constraints on F_0 Peak Timing in English" is a double study of Southern Standard British English, using natural uncontrolled data and a carefully designed corpus of read paragraphs. The timing of F_0 peaks is shown to be dependent on where the accent falls within a discourse unit: peaks were later in paragraph-initial position (equivalent in this study to discourse-topic-initial position) than in paragraph-internal position, and sentence-initial accents were in turn later than sentence-final ones. The authors conclude that topic-initiality exerts a strong rightward push on F_0 peaks, even causing them to occur outside the accented syllable. Predictably, discourse structure is also found to affect F_0 peak height, topic-initiality leading to higher peaks than topic-mediality.

In the Intonation and Technology section, Gösta Bruce, Marcus Filipson, Johan Frid, Björn Granström, Kjell Gustafson, Merle Horne, and David House's "Modelling of Swedish Text and Discourse Intonation in a Speech Synthesis Framework" provides an overview of an intonation model that was originally based on single-utterance laboratory speech. The model has now been extended to cover dialogues and multispeaker conversations, incorporating information on lexical semantics and discourse and textual structure. An important step in the research program is model-based resynthesis, whereby a synthetic F_0 contour is superimposed on the original utterance. The F_0 values for resynthesis are calculated on the basis of the symbolic utterance-level representation (pitch accents and boundary tones only) of the original. The differences between calculated and original F_0 values, such as overall trends and shifts up and down in F_0 range and height, are related to the analysis of the text in order to extract parameter values that can be fed into the text-to-speech implementation.

It is clear that this volume is aimed at readers who already have a basic knowledge of intonation and know what they are looking for. Since most of the chapters deal with highly specialized topics, each one is likely to be read in isolation. However,

an incentive to read more could have been provided, had each of the four sections been accompanied by a synopsis of its main themes and common threads.

Martine Grice is an assistant professor at Saarland University, Germany. Her main research interest is intonation theory, in particular, the structure of tonal representations. She has also developed schemes for the database annotation of tonal and junctural phenomena, both for Standard German (GToBI) and for a number of varieties of Italian (IToBI). Grice's address is Institute of Phonetics, FR.4.7, Saarland University, P.O. Box 151150, D-66041 Saarbrücken, Germany; e-mail: mgrice@coli.uni-sb.de.

Polysemy: Theoretical and Computational Approaches

Yael Ravin and Claudia Leacock (editors)

(IBM T. J. Watson Research Center and Educational Testing Services)

New York: Oxford University Press,
2000, xi+227 pp; hardbound, ISBN
0-19-823842-8, \$74.00, £45.00;
paperbound, ISBN 0-19-925086-3,
\$21.95, £14.99

Reviewed by

Jean Véronis

Université de Provence, Aix-en-Provence

As the editors of this volume remind us, polysemy has been a vexing issue for the understanding of language since antiquity.¹ For half a century, it has been a major bottleneck for natural language processing. It contributed to the failure of early machine translation research (remember Bar-Hillel's famous *pen* and *box* example) and is still plaguing most natural language processing and information retrieval applications. A recent issue of this journal described the state of the art in automatic sense disambiguation (Ide and Véronis 1998), and Senseval system competitions have revealed the immense difficulty of the task (<http://www.sle.sharp.co.uk/senseval2>). However, no significant progress can be made on the computational aspects of polysemy without serious advances in theoretical issues. At the same time, theoretical work can be fostered by computational results and problems, and language-processing applications can provide a unique test bed for theories. It was therefore an excellent idea to gather both theoretical and applied contributions in the same book.

Yael Ravin and Claudia Leacock are well-known names to those who work on the theoretical and computational aspects of word meaning. In this volume, they bring together a collection of essays from leading researchers in the field. As far as I can tell, these essays are not reprints or expanded versions of conference papers, as is often the case for edited works; instead, they seem to have been specially commissioned for the purposes of this book, which makes it even more exciting to examine.

The book is composed of 11 chapters. It is not formally divided into parts, but chapters dealing more specifically with the computational aspects of polysemy are grouped together at the end (and constitute about one-third of the volume).

Chapter 1 is an overview written by the volume editors. Yael Ravin and Claudia Leacock survey the main theories of meaning and their treatment of polysemy. These include the classical Aristotelian approach revived by Katz and Fodor (1963); Rosch's (1977) prototypical approach, which has its roots in Wittgenstein's *Philosophical Investigations* (1953); and the relational approach recently exemplified by WordNet (Fellbaum 1998), which (although the authors do not mention it) can be traced back to Peirce's (1931–1958) and Selz's (1913, 1922) graphs and which gained popularity with Quillian's (1968) semantic networks. In the course of this overview, Ravin and Leacock put the individual chapters into perspective by relating them to the various theories.

¹ The editors, citing Robins (1967), attribute the first observations of the "complex relations between meanings and words" to the Stoics, but reflection on polysemy can be traced back at least to Aristotle.

In Chapter 2, "Aspects of the Micro-Structure of Word Meanings," D. Alan Cruse addresses the issue of the extreme context-sensitivity of word meaning, which can result in an almost infinite subdivision of senses. However, Cruse believes that there are "regions of higher semantic density" within this extreme variability, which he calls sense-nodes, "lumps of meaning with greater or lesser stability under contextual change." As Cruse admits, this is only a metaphor, and as such, may not be highly useful to the researcher. In the rest of the chapter, Cruse attempts to build a typology of these nodes, listing their properties and providing tests to detect them. The tests (e.g., the zeugma effect in sentences such as *John and his driving license expired yesterday*) are not entirely new (e.g., Quine 1960; Cruse 1986; Geeraerts 1993), but are integrated here into a coherent framework that places context-dependency at the very heart of the theory.

Chapter 3 by Christiane Fellbaum is devoted to "autotroponymy." This term requires a two-step explanation. *Troponyms* are verb hyponyms, referring to specific manners of performing actions denoted by other verbs. For example, in English, *stammer*, *babble*, *whisper*, and *shout* are troponyms of *talking*. *Autotroponymy* is a special case that occurs when the verbs linked by this relation share the same form, as in *The children behaved / The children behaved well*. The author explains autotroponymy in terms of conflation of a meaning component not expressed on the surface. For example, in *The children behaved*, the verb includes a hidden adverbial (*well / satisfactorily / appropriately*). Fellbaum gives a typology of autotroponyms that is based on the nature of the conflated element (noun, adjective, adverbial), and she discusses their syntactic and semantic properties in detail.

In Chapter 4, "Lexical Shadowing and Argument Closure," James Pustejovsky explores verbs such as *butter*, which block the expression of a generic argument, as in **Mary buttered her bread with butter*, while allowing for a specific one, as in *Mary buttered her bread with expensive butter from Wisconsin* (see Levin 1993), and verbs such as *risk*, which can occur in contradictory contexts with roughly the same meaning, as in *Mary risked death to save her son / Mary risked her life to save her son* (see Fillmore and Atkins 1995). Pustejovsky introduces the concept of "lexical shadowing," which he defines as "the relation between an argument and the underlying semantic expression, which blocks its syntactic projection in the syntax." For example, the underlying semantics of the verb *butter* "shadows" the expression of the substance that is spread and allows only for specialization of the shadowed argument. For verbs such as *risk*, the shadowing is of a different type: it is the expression of one argument that shadows the expression of another, in a strictly complementary fashion. Pustejovsky explains these cases of argument optionality or complementarity in the framework of the Generative Lexicon (Pustejovsky 1995) and its various devices, among which "coercion" plays a central role.

Chapter 5, by Charles Fillmore and Sue Atkins, is a case study in lexicography. They analyze the sense divisions and definitions of the verb *crawl* in various dictionaries and compare them with corpus evidence from the British National Corpus. It is well known that dictionaries exhibit large discrepancies, and although they claim to be based on the analysis of corpus data, many sense distinctions that show up in a corpus are not reflected in dictionary entries. This is not entirely unexpected, since after all, no dictionary claims exhaustive coverage of a language, and some selection must be made by the lexicographer. This is even an explicit goal in four of the six dictionaries examined here, which are learners' dictionaries that attempt to illustrate the "core" uses of words for learners of English. It is striking, however, to see the extent to which lexicographers differ regarding their choices and assessment of what constitutes an important meaning a learner should acquire. Fillmore and Atkins are

perfectly right in noting that lexicographers lack objective criteria for sense division and information extraction from corpora. The FrameNet project they describe in an appendix (see <http://www.icsi.berkeley.edu/~framenet/>) is an attempt to achieve a systematic understanding and description of the meanings of lexical items and grammatical constructions by looking at a large number of attested examples, sorting them according to the conceptual structures (semantic “frames”) that underlie their meanings, and describing the associated information in terms of semantic roles, phrase types, and grammatical functions. The numerous observations regarding sense connections in the corpus examples result in a network-like organization of meanings, which can be used in both monolingual and bilingual lexicography. The last section of the chapter illustrates this possibility using the verb *ramper*, the French equivalent of *to crawl*.

Chapter 6, “‘The Garden Swarms with Bees’ and the Fallacy of ‘Argument Alternation’” by David Dowty, comes back to the argument problem already tackled by Fellbaum and Pustejovsky in their respective chapters and proposes syntactic structures as an explanatory principle for alternations in meaning. The author is concerned with agent / location alternations such as *Bees swarm in the garden* / *The garden swarms with bees*. He departs from the usual point of view that such pairs express the same meaning and differ only in syntactic form. Using the large set of examples in Salkoff (1983), Dowty groups verbs that participate in such alternations into five semantic classes and then shows that the two forms exhibit many semantic differences related to the informational structure of the sentence. The locative-subject form makes the location the topic of discourse, with the predicate ascribing an abstract property to the location. Some tests show the difference in meaning. For example, the *with*-phrase object must be semantically “unquantified” in the locative-subject form (compare *A roach crawled on the wall* / **The wall crawled with a roach*), the locative-subject form is more suited to metaphor than the agent-subject form, and so forth.

Chapter 7 by Cliff Goddard outlines Wierzbicka’s “natural semantic metalanguage” (NSM) approach to semantic analysis (Wierzbicka 1996, etc.), which is based on the idea that every language possesses a core of undefinable words (“semantic primes”). Complex expressions (words or grammatical constructions) can be described by means of explanatory reductive paraphrases composed of combinations of semantic primes. This “definitional” framework provides a diagnosis technique for detecting polysemy. For any given word, one can first assume that it has a single meaning and try to state it in a reductive paraphrase. If this turns out to be impossible and several paraphrases are needed to describe the word’s range of uses, then the word has distinct meanings. For example, there is no single paraphrase in terms of primes that could predict the range of uses of the French word *fille*, meaning both *daughter* and *girl*, and therefore the word must be split into two distinct meanings. Using this test, Goddard shows that dictionaries very often posit false and unnecessary polysemy, and occasionally false monosemy. He also shows how the technique can be used on grammatical constructions, and he applies it in detail to *have a VP* expressions (*have a stroll*, *have a chat*, etc.). The chapter ends with a discussion of how aspects of figurative language can be handled within this framework.

In Chapter 8, “Lexical Representations for Sentence Processing,” George Miller and Claudia Leacock raise the following question: “Why isn’t a dictionary a good theory of the lexical component of language?” They share Fillmore and Atkins’s dissatisfaction about dictionary making. For them, the main shortcoming of dictionaries is their lack of contextual information that would enable a user to make the correct association between senses and actual contexts. In their introduction, they give a convincing example from previous experiments. Schoolchildren given dictionary definitions of English

words produced sentences such as *Our family erodes a lot*, which sounds bizarre until you read the definition of *erode*: ‘eat out, eat away’. According to Miller and Leacock, what is missing from dictionaries is a satisfactory treatment of the lexical aspects of sentence processing. The rest of the chapter is devoted to a discussion of the two types of context that can be used to associate a given context with a particular word sense: local context (the immediate neighbors of the word under focus) and topical context (the general topic or domain of the text or conversation). The authors show that local context cues are very precise when they occur, but often simply do not occur. On the other hand, topical context is very efficient in helping discriminate between homographs, but not very helpful for identifying the different senses of a polysemous word. Miller and Leacock consider the combination of the two sources to be a major avenue of research.

Mark Stevenson and Yorick Wilks tackle this issue in Chapter 9, “Large Vocabulary Word Sense Disambiguation,” in which they propose a methodology for combining several knowledge sources into a word sense disambiguation system. Their first source of information is syntactic in nature and is provided by the Brill part-of-speech tagger. The semantic information present in the local context is then used in two ways. The overlap between *Longman Dictionary of Contemporary English* definitions and the local context is computed by means of an improved version of Cowie, Guthrie, and Guthrie’s (1992) simulated-annealing technique, and selectional restrictions are resolved by means of *LDOCE* semantic classes. The larger context is handled with techniques that map it to the subject categories provided by *LDOCE* for each sense (“pragmatic codes”). The efficiency of each of these modules taken separately ranges from 44 to 79 percent, but Stevenson and Wilks show that using machine learning techniques, the modules can be combined in an efficient way to produce 90 percent correct disambiguation, which is quite high for an unrestricted vocabulary system.

In Chapter 10, “Polysemy in a Broad-Coverage Natural Language Processing System,” William Dolan, Lucy Vanderwende, and Steven Richardson describe the approach to polysemy processing taken in the MS-NLP broad-coverage natural language understanding system. The core of their system is MindNet, a network-structured computational lexicon extracted from machine-readable dictionaries (MRD) augmented with corpus information. MindNet uses the same general approach as the MRD-based spreading activation networks proposed by Véronis and Ide (1990), although in a much more sophisticated version including labeled connections, backward links, weighted paths, and so on. The authors depart from most computational approaches to polysemy in that they believe that word meaning is “inherently flexible,” that making predefined inventories of discrete senses is unsuitable for broad-coverage applications, and that no sharp boundaries should be drawn between senses. Their approach is reminiscent of Cruse’s, presented earlier in this book. For these authors, “understanding” is no more than identifying an activation pattern in the network.

In previous publications, Hinrich Schütze held a position similar to Dolan, Vanderwende, and Richardson’s with respect to predefined sense inventories. For Schütze, many problems require *discrimination* among senses but do not require explicit sense *labeling*, and the techniques he has proposed extract the sense divisions from the corpus itself (see Schütze 1998): a sense is a group of contextually similar occurrences of a word. This approach is almost the opposite of Goddard’s. In Chapter 11, Schütze looks at word sense disambiguation from the perspective of connectionism. After surveying some of the literature on disambiguation, he presents an algorithm that has grown out of two major concerns in connectionist research: psychological plausibility and large-scale applicability. He describes an application to information retrieval

that demonstrates that his algorithm can be applied to very large text collections (500 megabytes of text from the *Wall Street Journal*).

The most noticeable feature of this book is probably its wide range of contributors and the broad scope of the topics it encompasses. As the title implies, it addresses both theoretical and computational aspects of polysemy, and within these two areas, very different research trends are pursued. The book gives a very good overall picture of current issues in polysemy and of the diverse ways of approaching the topic. It should therefore hold an important place on the shelves of any researcher in the fields of lexical semantics and word sense disambiguation, and will certainly be valued by many of our graduate students.

The wide-angle snapshot offered by this book also reveals a very striking fact about current lexical semantics. Apart from one chapter, all theoretical discussions are supported solely by invented examples. Lexical semantics, and probably semantics in general, has not yet made the paradigm shift that has occurred or is occurring in other branches of linguistics, such as syntax, where empirical evidence now replaces intuition as the normal body of data to be studied. Another recent book (Sampson 2001) quite brilliantly shows how the lack of objective evidence has been misleading linguistic research for decades and has placed the discipline on the fringe of modern science. The lack of objective evidence is probably even more dangerous in semantics than in other areas of linguistics. The extreme flimsiness of introspection-based tests is acknowledged by lexical semanticists themselves—for instance, how much agreement would there be on whether or not a given coordination is a zeugma?—and such tests make it almost impossible for semantics to satisfy the minimal requirement that science has demanded since Karl Popper, that of refutability.

Interestingly enough, the one chapter that does use corpus examples (Chapter 5 by Fillmore and Atkins) pertains to lexicography. Lexicographers indeed have a long tradition of examining objective evidence, which computer tools and electronic corpora have made it possible to systematize. However, several chapters (Chapter 5 by Fillmore and Atkins, Chapter 7 by Goddard, Chapter 8 by Miller and Leacock) express their dissatisfaction with current dictionaries, on the grounds that they lack theoretical criteria to back their organization. It is also worth noting that the only computational approaches to word sense disambiguation able to claim some minimal degree of efficiency are linguistically blind ones (like those reported in this book), as if an insurmountable gap existed between theories and applications. A paradigm shift in lexical semantics is therefore not just a scientific necessity; it is also a practical one. I am convinced that no major breakthrough in language-processing applications and lexicography can be made until theories of meaning are based on the observation of real data.

References

- Cowie, Jim, Joe A. Guthrie, and Louise Guthrie. 1992. Lexical disambiguation using simulated annealing. In *Proceedings of the 14th International Conference on Computational Linguistics (COLING'92)*, 23–28 August, Nantes, France, vol. 1, pages 359–365.
- Cruse, D. Alan. 1986. *Lexical Semantics*. Cambridge University Press, Cambridge, UK.
- Fellbaum, Christiane. 1998. *WordNet: An Electronic Lexical Database*. MIT Press, Cambridge, MA.
- Fillmore, Charles J. and Beryl T. Atkins. 1995. Toward a frame-based lexicon: The semantics of RISK and its neighbors. In Adrienne Lehrer and Eve Feder Kittay, editors, *Frames, Fields, and Contrasts*. Lawrence Erlbaum Associates, Hillsdale, NJ, pages 75–102.
- Geeraerts, Dirk. 1993. Vagueness's puzzles, polysemy's vagaries. *Cognitive Linguistics*, 4(3):223–272.
- Ide, Nancy M. and Jean Véronis. 1998. Introduction to the special issue on word sense disambiguation: The state of the art. *Computational Linguistics*, 24(1):1–40.

- Katz, Jerrold J. and Jerry A. Fodor. 1963. The structure of a semantic theory. *Language*, 39:170–210.
- Levin, Beth. 1993. *English Word Classes and Alternations: A Preliminary Investigation*. University of Chicago Press, Chicago.
- Peirce, Charles Sanders. 1931–1958. *Collected Papers of C. S. Peirce*, edited by C. Hartshorne, P. Weiss, and A. Burks, 8 vols., Harvard University Press, Cambridge, MA.
- Pustejovsky, James. 1995. *The Generative Lexicon*. MIT Press, Cambridge, MA.
- Quillian, M. Ross. 1968. Semantic memory. In Marvin Minsky, editor, *Semantic Information Processing*. MIT Press, Cambridge, MA, 227–270.
- Quine, Willard Van Orman. 1960. *Word and Object*. MIT Press, Cambridge, MA.
- Robins, Robert H. 1967. *A Short History of Linguistics*. Indiana University Press, Bloomington.
- Rosch, Eleanor. 1977. Human categorization. In N. Warren, editor, *Advances in Cross-Cultural Psychology*, vol. 7. Academic Press, London.
- Salkoff, Morris. 1983. Bees are swarming in the garden. *Language*, 59(2):288–346.
- Sampson, Geoffrey. 2001. *Empirical Linguistics*. Continuum, London.
- Schütze, Hinrich. 1998. Automatic word sense discrimination. *Computational Linguistics*, 24(1):97–124.
- Selz, Otto. 1913. *Über die Gesetze des geordneten Denkoerlaufs*. Spemann, Stuttgart.
- Selz, Otto. 1922. *Zur Psychologie des produktiven Denkens und des Irrtums*. Friedrich Cohen, Bonn.
- Véronis, Jean and Nancy Ide. 1990. Word sense disambiguation with very large neural networks extracted from machine readable dictionaries. In *Proceedings of the 13th International Conference on Computational Linguistics (COLING'90)*, Helsinki, Finland, vol. 2, pages 389–394.
- Wierzbicka, Anna. 1996. *Semantics: Primes and Universals*. Oxford University Press, Oxford.
- Wittgenstein, Ludwig. 1953. *Philosophical Investigations* (translated by G. E. M. Anscombe). Macmillan, New York.

Jean Véronis is a professor of linguistics and computer science at the Université de Provence in Aix-en-Provence, France, where he heads a research team specializing in French corpus linguistics. His academic interests include word sense disambiguation, computer lexicography, translation corpora and parallel text alignment, prosody, and speech synthesis. Véronis's address is Université de Provence, 29 av. Robert Schuman, 13621 Aix-en-Provence Cedex 1, France; e-mail: Jean.Veronis@up.univ-mrs.fr; URL: www.up.univ-mrs.fr/~veronis.

Abduction, Belief and Context in Dialogue: Studies in Computational Pragmatics

Harry Bunt and William Black (editors)
(Tilburg University and UMIST)

Amsterdam: John Benjamins (Natural
language processing series, edited by
Ruslan Mitkov, volume 1), 2000,
vi+471 pp; hardbound, ISBN
1-55619-794-2 and 90-272-4983-0, \$90.00

Reviewed by
Matthew Stone
Rutgers University

The problem of pragmatics arises as soon as we move beyond the linguistic analysis of an utterance and ask what the speaker meant by it. Now, speaker's meaning is a particular case of intention, and all intentions are complex mental attitudes that tie an agent's actions to its goals, its background beliefs, and its appraisal of the context in which it acts. Inferring the speaker's meaning, then, is a matter of recognizing how the speaker might have represented the utterance as linking up with the current discourse context and thereby furthering the goals of the conversation. This recognition of intention is reasoning to the best explanation, or abductive reasoning, as popularized in computational linguistics by Hobbs and colleagues (1993). So there they are: abduction, belief, and context—Harry Bunt and William Black's ABC of computational pragmatics—three concepts that rightly frame the diversity of current research in problems of pragmatic interpretation.

In this volume (*ABC* for short), Bunt and Black collect 15 chapters that grow, by and large, out of the ESPRIT project PLUS (Pragmatics-Based Language Understanding System). This coherence lends the volume strengths not often found in collections of research papers; *ABC* fits together to give a broad picture of computational pragmatics as an interdisciplinary enterprise in which a multitude of different investigations can be brought to bear constructively on a common project.

To start, *ABC* offers generally consistent terminology and perspective, laid out in a 150-page three-chapter overview of computational pragmatics in PLUS. As outlined by Bunt and Black in Chapter 1, dialogue system design in PLUS centered on the representations and inference required for recognizing the communicative intentions behind users' utterances. Jens Allwood's Communicative Activity Analysis (Chapter 2) provided the theoretical framework for this design, while Bunt's Dynamic Interpretation Theory (Chapter 3) bridged this dialogue theory and specialized models of meaning and context from computational linguistics and from computer science more generally.

Concretely, this approach is distinguished by a wide view of context, including social and physical dimensions as well as linguistic ones, and a wide view of agency, mandating considerations of ethics and trust in cooperation from the start. It is also guided by some more practical working assumptions:

- a focus on dialogues of information exchange in which domain reasoning (particularly reasoning about users' domain plans and domain-specific communication strategies) can be sharply circumscribed;

- an emphasis on the conventionality of moves in dialogue, against indirect speech acts and other particularized conversational implicatures;
- an eclectic use of mathematical tools for describing information states in dialogue—from computational logic, knowledge representation, and deductive databases as well as computational linguistics.

ABC also benefits from its coherence in achieving impressive coverage with minimal redundancy. In addition to the introductory material, chapters cover general problems in implementing interactive dialogue systems: context representation (Bunt), user modeling (Meyer), and architectures for system design (Sabah; Taylor and Waugh). Other chapters cover more specifically linguistic aspects of computational pragmatics: attention tracking (Carter), discourse structure (Redeker), and speech acts (Ramsay; Beun; Thijsse). Finally, connecting problems in computational linguistics and interaction are chapters on inference for dialogue understanding and generation, united by the theme of abduction (Neal): abductive generation (Oberlander and Lascarides), abductive context updating (Guessoum and Gallagher), and abductive interpretation (Hinkelman and Spackman). The chapters also illustrate a wide range of methodology, including not only the ubiquitous formal modeling but also corpus analysis (Redeker), Wizard-of-Oz studies (Beun), system building (Hinkelman and Spackman), and evaluation (Carter).

The coherence of *ABC* has the unfortunate side effect of offering little comparison with other ongoing work in computational pragmatics beyond the projects the book documents. *ABC*'s assumptions about dialogue are in fact rather controversial. Is information exchange simple? Research on cooperative response has made information exchange the classic test bed for modeling domain problem-solving and its ramifications for dialogue. Once a user makes clear his intention to *do something* with the information he gets, you may have a full task-oriented dialogue with all the structure and communicative action that entails. Can dialogue moves be purely conventional in a robust system? Or are disagreements, misconceptions, accommodation, even jokes, so common and so rich that a first-principles representation of communicative intention along Gricean lines cannot be avoided? And where are we to push for representations of context: database theory or models of uncertainty? With more and better comparison to alternative models, *ABC* would have done much better at conveying the difficulty, vitality, and diversity of research in computational pragmatics.

Even if *ABC* is hardly a source for all the latest ideas in dialogue,¹ ultimately, as always, it is the field itself that will make lasting comparisons. With its broad, consistent tutorial flavor, *ABC* deserves a place right behind Cohen, Morgan, and Pollack's *Intentions in Communication* (1990) as an accessible introduction to some classic ideas in the computational analysis of conversation.

References

Cohen, Philip R., Jerry Morgan, and Martha E. Pollack, editors. 1990. *Intentions in Communication*. MIT Press, Cambridge, MA.

Hobbs, Jerry, Mark Stickel, Douglas Appelt, and Paul Martin. 1993. Interpretation as abduction. *Artificial Intelligence*, 63:69–142.

¹ Indeed, *ABC* seems to have had the usual protracted publication. The PLUS project ran from 1990 to 1994. Within *ABC*, its other chapters are cited from 1998 through 2000; on the Web, chapter drafts show up as early as 1995.

Matthew Stone is assistant professor in the Computer Science Department and the Center for Cognitive Science at Rutgers, the State University of New Jersey. His research explores the role of representations of pragmatic interpretation in explaining human-human dialogue and constructing conversational systems. Stone's address is 110 Frelinghuysen Road, Piscataway, NJ 08854-8019; e-mail: mdstone@cs.rutgers.edu; URL: <http://www.cs.rutgers.edu/~mdstone>.

Relationships in the Organization of Knowledge

Carol A. Bean and Rebecca Green (editors)
(University of Tennessee and University of Maryland)

Dordrecht: Kluwer Academic
Publishers (Information science and
knowledge management series, edited
by J. Mackenzie Owen, volume 2), 2001,
ix+232 pp; hardbound, ISBN
0-7923-6813-4, \$97.00, £61.00, €90.00

Reviewed by
Gregory Grefenstette
Clairvoyance Corporation

This book is not for linguists. This book is not for computer scientists. This book is for undergraduate librarians.

Given the promising title, from which one might expect a comprehensive look at imposing order on knowledge, this is a little disappointing. The back-cover blurb insists that the book can be used to provide “guidance for relational tasks [that are] now taking on greater significance, as retrieval systems increasingly operate in automated modes and as retrieval systems cross linguistic, cultural, and disciplinary boundaries.” But going inside, one enters the closed-world community of library science for which the automation of the card catalogue is still big news.

The book is disappointing because so many strands of modern computational and linguistic work attacking the problem of organizing knowledge are missing. For example, from computer science, the fact that relationships have formed the core of database modeling ever since relational databases (Codd 1970) and the entity-relationship model (Chen 1976) goes unacknowledged. From artificial intelligence, missing is any of the work on ontologies for organizing knowledge, such as Penman’s upper model (Bateman et al. 1990) and CYC (Lenat et al. 1990), and work on organizing knowledge in expert systems. As for linguistics, there is no mention of the strain of research stemming from the idea of semantic markers (Katz and Fodor 1963) and the myriad subsequent efforts to exploit them to represent meaning despite their inherent limitations (Bolinger 1965; Eco 1976). All this practical and intellectual effort is ignored here. The content of this book is light years away from current practical concerns of knowledge representation (www.kr.org), knowledge management (www.cikm.org), or the relationships within domains of knowledge being explored with UML (www.uml.org) and XML (www.xml.org). Maybe these criticisms all derive from false hopes raised by the book’s title. A better title for this book would be *Relationships in the Organization of Library Books*.

The first part of the book is “Theoretical Background.” Rebecca Green opens the book by describing the wide variety of efforts to characterize relations from the bibliographic and documentary points of view. Barbara Tillet follows with a brief, low-level tutorial on how bibliographic items can be related to each other—for example, how to describe “works within works,” in which one book contains other books. Next, Stella Dextre Clarke describes the International Standards Organization (ISO) relations used in manually built thesauri (e.g., BT for broader term, NT for narrower term, and their multiple variants). Here, Dextre Clarke shows how the same word is described incompatibly in different thesauri; she also shows in a table a number of ways in which the

RT (related term) relationship actually covers a wide variety of relations (e.g., “whole-part,” “action and its patient,” “causal dependence,” “concept and its opposite,” and 11 other relations with examples). Jessica Milstead, in her chapter about standards for relations between subject-indexing terms, laments how much latitude is possible in interpreting relations mentioned in thesaural standards. Their interpretations, she says, may be reasonably evident to a human manipulating the thesaurus, but she warns that “[i]t is when the thesaurus sits behind the interface as a tool to be applied [automatically] by the system that more precise specification of equivalence relations is required”.

Michele Hudon describes the drawbacks of multilingual thesauri, so that, for example, even though words may pass as translations, restriction of meanings makes thesaural relations between near-translations quickly diverge. For example, translations of the hyponyms of the English word *education* are not hyponyms under the French word *éducation*, which, though a near-translation, has a more restricted meaning than the English word. Some of the translations would have to be hyponyms of the French word *enseignement*. At the end of her contribution, Hudon refers to work on EuroWordNet (Vossen 1998) that has been addressing this problem.

Olivier Bodenreider and Carol A. Bean, one of the editors, talk about some of the problems involved in integrating different thesauri into one coherent organization, giving details about the structure of shared medical vocabularies of the Unified Medical Language System (UMLS). This chapter contains many pointers to medical terminological resources on the Web, to which one could add the International Medical Informatics Association Working Group 6 on Medical Concept Representation (www.mayo.edu/imia-wg6/).

Clare Beghtol offers a philosophical musing on structure and meaning, following a historical trajectory from Aristotle through to the Dewey decimal system and social psychology. The editors then come back with a chapter on relevance relations. This chapter might be at home in a book about information retrieval (which is adequately cited here), but seems out of sync with the rest of the book, which hovers around classic cataloguing and thesaurus use. In this chapter, Bean and Green present some of their previous research that attempted to typify the relations between a topical guide to scriptures and the passages cited as relevant to the topics in the guide. The research found that “topical relevance relations include a wide variety of relationships, only some of which, perhaps only a relatively small proportion, are matching relations” (page 127). Considering that matching relations are what people use now to browse the Web, this conclusion is interesting, as is the rest of this chapter’s discussion about what relevance means.

The second part of the book, titled “Systems,” contains chapters reviewing Library of Congress subject headings, the structure of the Art and Architecture Thesaurus (www.getty.edu/research/tools/vocabulary/aat/), Medical Subject Headings (MeSH, www.nlm.nih.gov/mesh), a multicultural and multilingual thesaurus from India, Colon Classification (an attribute-value-type library classification system developed by S. R. Ranganathan in the early 1950s), and the Dewey decimal system. This part of the book might be useful to a computational linguist who needs a quick reference to one of the classic classification systems devised for organizing library stocks.

References

Bateman, John, Robert Kasper, Johanna Moore, and Richard Whitney. 1990. A general organization of knowledge for

natural language processing: The Penman upper model. Technical Report, Information Sciences Institute, Marina del Rey, CA.

- Bolinger, Dwight. 1965. The atomization of meaning. *Language*, 41(4):555–573.
- Chen, Peter Pin-Shan. 1976. The entity-relationship model: Towards a unified view of data. *ACM Transactions on Database Systems*, 1(1):9–36.
- Codd, E. F. 1970. A relational model for large shared data banks. *Communications of the ACM*, 13(6):377–387.
- Eco, Umberto. 1976. *A Theory of Semiotics*. Indiana University Press, Bloomington.
- Katz, Jerrold J. and Jerry A. Fodor. 1963. The structure of a semantic theory. *Language*, 39(2):170–210.
- Lenat, Douglas B., Ramanathan V. Guha, Karen Pittman, Dexter Pratt, and Mary Shepherd. 1990. CYC: Toward programs with common sense. *Communications of the ACM*, 33(8):31–49.
- Vossen, Piek. 1998. *EuroWordNet*. Kluwer, Dordrecht.

Gregory Grefenstette, Principal Research Scientist at Clairvoyance Corporation, is the author of *Explorations in Automatic Thesaurus Discovery* (Kluwer Academic Publishers, 1994) and editor of *Cross-Language Information Retrieval* (Kluwer Academic Publishers, 1998). He has reviewed manuscripts for many natural language processing, artificial intelligence, and computational linguistics journals and conferences. He is currently working on very large lexicons. Grefenstette's address is Clairvoyance Corporation, 5301 Fifth Avenue, Pittsburgh, PA 15232; e-mail: g.grefenstette@Clairvoyancecorp.com.