

Book Reviews

The Mathematics of Language

Marcus Kracht

(University of California, Los Angeles)

Berlin: Mouton de Gruyter (Studies in generative grammar, volume 63), 2003, xvi+589 pp; hardbound, ISBN 3-11-017620-3, \$127.00, €98.00

Reviewed by

Hans-Jörg Tiede

Illinois Wesleyan University

Mathematical linguistics is concerned with the study of mathematical properties of natural languages and linguistic theories. Since the mathematical properties of interest to mathematical linguists are usually from theoretical computer science (complexity classes, language hierarchies, formal learnability), mathematical linguistics can be considered to be an area of theoretical computational linguistics. However, since statistical methods are rarely used in mathematical linguistics, its relationship to current practices in computational linguistics is somewhat limited.

While the introduction of logic in linguistic research originally came from semantics, this line of work did not really use sophisticated metaresults. One of the main developments in mathematical linguistics in the last decade has been the introduction of sophisticated logical methods to the study of natural language syntax, for instance, the use of cut elimination and interpolation theorems in categorial grammar or Büchi's theorem relating finite automata and monadic second-order logic in model-theoretic syntax. The book under review is written by one of the main contributors to the logical turn in mathematical linguistics, and so it is not surprising that this is its main focus.

The contents of the book are as follows: Chapter 1, "Fundamental Structures," gives a concise introduction to the mathematical background needed for the rest of the book. It does not give an introduction to logic at this point, which is introduced as needed in the rest of the book. Chapter 2, "Context Free Languages," starts with the regular languages and then discusses normal forms, parsing, and ambiguity. It concludes with a proof of Parikh's theorem, which states that every context-free language is semilinear, and a discussion of non-context-free phenomena in natural languages. Chapter 3, "Categorial Grammar and Formal Semantics," contains, in addition to an introduction to the λ -calculus and combinators, an introduction to the Lambek calculus, culminating in a complete proof of Pentus's theorem, which states that Lambek grammars are context-free. Chapter 4, "Semantics," introduces algebraic tools for the study of natural language semantics. Chapter 5, "PTIME Languages," discusses extensions of context-free grammars, including tree-adjoining grammars, indexed grammars, and literal movement grammars. It includes a discussion of the class of mildly context-sensitive languages, which are widely held to be fairly good approximations to the complexity of natural languages. The last chapter, chapter 6, "The Model Theory of Linguistic Structure", is an introduction to us-

ing the approach of descriptive complexity theory to define conditions on strings and trees, as well as phonological representations. The chapter concludes with applications of the logical tools to grammar formalisms, including GPSG, HPSG, and GB.

As can be seen from the contents, the main focus of this book is on syntax, specifically the application of formal language theory and logic to natural languages. At this point, the classical introduction to this area is still Partee, ter Meulen, and Wall (1990). Kracht's book is significantly more advanced; in fact, a good working knowledge of Partee et al. is a prerequisite to Kracht's book. The *Handbook of Logic and Language* (van Benthem and ter Meulen 1997), which contains advanced material in this area, is much more focused on semantics than Kracht's book. Compared to introductory (e.g., Hopcroft and Ullman 1979) or advanced (e.g., Martín-Vide, Mitrana, and Păun 2004) books on formal language theory, Kracht's book emphasizes those aspects of formal language theory that are relevant to the study of natural languages, whereas the former do not. Thus, Kracht's book gives a uniform introduction, which currently does not exist at this level, to an important area of mathematical linguistics. Its main use will be in advanced graduate courses and for researchers interested in learning about mathematical linguistics.

The book stems from lecture notes that the author produced for a number of classes in this area; however, on the continuum that ranges from textbooks to research monographs, this book is located somewhere in the middle. For instance, from the point of view of a textbook, the importance that monadic second-order logic currently plays in mathematical linguistics would have warranted spending a larger part of chapter 6 on it rather than on quantified modal logic, which is a somewhat idiosyncratic choice. Such choices occur at a few other places in the book; however, they are balanced by the almost encyclopedic overview of formal grammars and important results about them that cannot currently be found in one volume. Another strength of this book is that it introduces logical tools incrementally together with the application for which they are needed. This reinforces the central role that logic plays in mathematical linguistics and makes it possible to read the book as an introduction to applied logic.

There are some minor inconsistencies. For instance, indexed grammars are introduced in the chapter on PTIME languages, and it is claimed that languages generated by indexed grammars can be parsed in PTIME, even though the recognition problem for indexed grammars is known to be NP-complete (Rounds 1973).

The chapter on semantics is a little too brief. Given the length of the book, it is understandable that no more space could be dedicated to it; however, it might have been more productive to replace that material with a more detailed description of formal models of GB, since this is an area in which Kracht has obtained many important results. Computational linguists will probably feel that parsing is not discussed in enough detail. While chart parsing is discussed for context-free grammars, a large number of papers on parsing of mildly context-sensitive grammars can be found in conference proceedings, but there is no uniform introduction to this area at this time.

While the book contains many exercises of different degrees of difficulty, I couldn't find any open research problems. This is unfortunate given that most readers of this book will likely be researchers and Ph.D. students.

Overall this book is an excellent introduction to advanced topics in mathematical linguistics that, given its advanced nature, requires a significant amount of mathematical maturity.

References

- Hopcroft, John E. and Jeffrey D. Ullman. 1979. *Introduction to Automata Theory, Languages, and Computation*. Addison-Wesley.
- Martín-Vide, Carlos, Victor Mitrană, and Gheorghe Păun, editors. 2004. *Formal Languages and Applications*. Springer.
- Partee, Barbara, Alice ter Meulen, and Robert E. Wall. 1990. *Mathematical Methods in Linguistics*. Kluwer.
- Rounds, William C. 1973. Complexity of recognition in intermediate-level languages. In *14th Annual IEEE Symposium on Switching and Automata Theory (University of Iowa, Iowa City, Iowa, 1973)*, pages 145–158. IEEE Computer Society, Northridge, CA.
- van Benthem, Johan and Alice ter Meulen, editors. 1997. *Handbook of Logic and Language*. MIT Press.

Hans-Jörg Tiede is an assistant professor of computer science at Illinois Wesleyan University. His research interest is the application of logic to mathematical and computational linguistics. Tiede's address is Department of Mathematics and Computer Science, Illinois Wesleyan University, P.O. Box 2900, Bloomington, IL 61702-2900; e-mail: htiede@iwu.edu.

Recent Advances in Example-Based Machine Translation

Michael Carl and Andy Way (editors)

(Universität des Saarlandes and Dublin City University)

Dordrecht: Kluwer Academic
Publishers (Text, speech and language
technology series, edited by Nancy Ide
and Jean Véronis, volume 21), 2003,
xxxi+482 pp; hardbound, ISBN
1-4020-1400-7, \$173.00, £115.00, €180.00

Reviewed by

Walter Daelemans

University of Antwerp

This book, an outcome of a 2001 workshop on example-based machine translation (EBMT) in Santiago de Compostela, very appropriately starts with a preface by professor Makoto Nagao in which he explains how the limits of rule-based machine translation (MT) led him to propose his principle of translation by analogy in 1981 (published as Nagao 1984). His idea, inspired by second-language learning methodology, is convincing and elegant. We learn a language not by deep linguistic analysis and rule application, but by rote learning and analogy-based generalization. Starting from simple sentences and their translations, variations of these sentences can also be translated by using similarity-based reasoning. If the Japanese translation of a sentence such as *A man eats vegetables* is memorized (implicitly including changes in word order, morphology, translation selection of *eat*, etc.), then a sentence such as *He eats potatoes* can be translated correctly by analogy to this memorized case if the semantic relation between *he* and *a man* and between *vegetables* and *potatoes* can be determined (e.g., using a thesaurus). Such a method bypasses deep linguistic analysis by leaving transfer rules implicit in an aligned bilingual corpus.

Michael Carl and Andy Way have done the field of MT and that of natural language processing in general a big favor by producing this collection. This is the first (fat) book-sized collection of current research in this fascinating area of MT, and it contains a number of chapters that will bring anyone interested quickly up to speed in this research area. In order to differentiate EBMT from statistical MT and rule-based MT, it is tempting to describe it as a hybrid technique taking a middle ground between rule-based methods (linguistic knowledge used in the representation of translation units) and data-oriented methods (learning from bilingual parallel corpora, similarity-based reasoning). But in reality, EBMT appears in this book as a widely varying bunch of somehow related approaches: a concept that can best be described in an example-based way. In their introduction, Carl and Way acknowledge this lack of an analytical definition but see no harm in it and compare the situation with other “scientific realities” such as artificial intelligence. The volume is divided into four parts, the first being about historical and technical foundations, and the other three describing current research.

Part I contains a wealth of background information. Chapter 1, by Harold Somers (“An Overview of EBMT”) traces the history of the idea and dives into the underlying problems to be solved in applying it: collecting and aligning parallel corpora, deciding on the grain size and linguistic representation of the examples (sentence level or below, superficial or deep?), the number and selection of examples needed, storage, indexing

and matching of examples, and issues involved in using the best matching examples to produce grammatical output (adaptation of the retrieved knowledge). The remainder of the chapter discusses different flavors, extensions, and uses of EBMT and describes how EBMT has been incorporated in multiengine systems, in parallel with rule-based and statistical systems. The anecdotal, impressionistic, and limited evaluation of EBMT available in the literature according to Somers is surprising. Especially given the trend toward multiengine systems, it is to be hoped that a thorough comparative evaluation of EBMT will become available soon. Somers mentions corpus coverage, simple extension, linguistic theory-freeness, and easy development from aligned corpora as the main advantages of EBMT and scalability and dealing with some translation problems as its main problems. He sees it not as a rival to rule-based MT but as an enhancement of it. Given the results of, for example, Sumita in a later chapter of the book, I think he gives up on “pure” EBMT too easily.

Chapter 2 by Davide Torcato and Fred Popowich (“What is Example-Based Machine Translation?”) attempts to single out what sets EBMT apart from other approaches. The authors argue that the linguistic information being used by an MT system (rather than how it is represented or acquired) should be the basis for classification. On this basis, they go to great length contending that most of the properties of EBMT (use of parallel corpus, holistic approach, etc.) are not really distinguishing properties and that only the translation-by-analogy principle stands out as truly example-based.

In chapter 3 (“EBMT in a Controlled Environment”), Reinhard Schärer, Andy Way, and Michael Carl see EBMT (in the guise of phrasal lexicons) as a way of enhancing translation memories (TM). TM is arguably the commercially most successful approach to MT. The idea is probably as old as EBMT (I heard it being mentioned in the very first workshop I ever attended in 1984 in Davos). In TM, a translator has access to previous translations of sentences that are retrieved by (fuzzy) matching. Such an approach will have high precision but very low recall. Schärer et al. demonstrate that by extending the approach to the subsentential level (phrases), recall can be improved. The main contributions of this paper are pointing to the use EBMT might have in improving the usefulness and coverage of TM systems (although that point has been made by many others as well), analyzing how this could be done in the context of controlled language TM, and what the benefits and limitations could be. Unfortunately it remains a position paper more than a proof of concept.

The foundational part of the book ends with an interesting chapter by Bróna Collins and Harold Somers (“EBMT Seen as Case-Based Reasoning”) in which MT is analyzed in the framework of case-based reasoning (CBR), a well-known reasoning approach in AI. The goal is to investigate whether other applications of CBR can contribute to its application to MT. Unfortunately, the chapter doesn’t go very deep into analyzing the relations and differences between different kinds of “lazy learning,” of which CBR is an instance (memory-based, instance-based, and exemplar-based reasoning and learning methods have all been proposed in the AI literature). I would refer people interested in that topic rather to Aha (1997) or Kasif et al. (1998). Nevertheless, the explicit analogy between the two analogical approaches in this chapter is extensive and insightful and may inspire new research on extensions and variations of EBMT.

EBMT systems differ in whether they acquire the knowledge implicit in the bilingual corpus at runtime or off-line in advance. Runtime approaches, of which a few are presented in part II of the book, extract their translation units during translation from a sentence-aligned bilingual corpus. Chapter 5 (“Formalizing Translation Memory”) by Emmanuel Planas and Osamu Furuse addresses in a sense the same issues as chapter 3, trying to extend translation memory approaches to EBMT. The way

this is done in their “shallow translation” approach is to add lemmas, POS tags, and string-edit distance on multiple levels of sentences for matching in TM. Preliminary evaluation comparing the matching approach to a standard TM-matching algorithm shows that especially at high similarity thresholds, their approach retrieves more useful cases. With the contribution of Eiichiro Sumita (“An Example-Based Machine Translation System using DP-Matching between Word Sequences”), we get near to the pure EBMT approach, generalizing examples on the fly and not using any syntactic parsing or bilingual treebanks. On a reasonably sized corpus (200,000 sentences), the approach covers about 90% of the 500 test sentences, with about 80% acceptable translations, taking 10 seconds per sentence on average. This is hopeful news for the pure approach, although many problems remain, especially regarding handling of sparse data, long sentences, and context dependency in the approach. Francis Bond and Satoshi Shirai (“A Hybrid Rule and Example-Based Method for Machine Translation”) contribute an approach to combine the strengths of rule-based and example-based approaches, and Tantely Andriamanankasina, Kenji Araki, and Koji Tochinnai (“EBMT of POS-Tagged Sentences by Recursive Division via Inductive Learning”) apply the CBR approach to subsententially aligned examples. Both approaches claim promising results.

Parts III and IV of the book concern approaches to and systems of EBMT that do not acquire their knowledge from the bilingual corpus dynamically and on-line but rather do it off-line, either as extracted translation templates (part III) or as something that starts to resemble structural transfer rules (part IV). Notice that from the point of view of a machine-learning interpretation of example-based reasoning and learning, something bizarre is happening here. The techniques described in these last two parts of the book become increasingly less example-based. It is not because rules are learned from examples that an approach becomes example-based; the crucial aspect is that the examples themselves are used in reasoning, not generalizations extracted on the basis of them. The fundamental difference between rule-based and example-based approaches is that the former, because of the nature of rules, have to abstract from low-frequency and untypical examples in order to formulate compact rules, whereas the latter keep all information, exceptions, and noise included. This does not imply that the work described in these papers is any less interesting, of course, but I would not necessarily call it EBMT. Ilyas Cicekli and H. Altay Güvenir (“Learning Translation Templates from Bilingual Translation Examples”) show how translation templates can be learned by means of a language-independent method for generalizing exemplars based on similarities and differences. Ralf Brown (“Clustered Transfer Rule Induction for Example-Based Translation”), meanwhile, adopts a similar approach to that used in the previous chapter to learn translation templates and adds to that a bottom-up agglomerative clustering method for both words and replacement rules. He shows that clustering and rule induction each outperform simple string matching and that the combination outperforms both. In chapter 11 (“Translation Patterns, Linguistic Knowledge and Complexity in an Approach to EBMT”), Kevin McTait discusses a template extraction based on similarity in distributions of strings in source and target language sentences and fails to improve its accuracy significantly by adding morphological analysis and POS tagging. Finishing Part III, Michael Carl (“Inducing Translation Grammars from Bracketed Alignments”) presents a system that extracts lexical transfer rules and translation templates from a tagged and bracketed corpus, thereby effectively moving from example-based reasoning to grammar induction. Interestingly, the induced grammar has the desirable properties of homomorphy, invertibility, and compositionality.

Part IV moves even further away from example-based approaches in requiring the extraction of translation grammars from structured representations (bilingual treebanks). Kaoru Yamamoto and Yuji Matsumoto, in chapter 13 (“Extracting Translation

Knowledge from Parallel Corpora”), report on two successful studies on extracting translation knowledge from parallel automatically annotated corpora, with robust results even with the unavoidable annotation errors. They also show that chunk boundaries, especially, provide useful information for translation and that dependency relations are crucial for longer phrasal translation pairs. Hideo Watanabe, Sadao Kurohashi, and Eiji Aramaki (“Finding Translation Patterns from Paired Source and Target Dependency Structures”) follow a similar approach but add a method for extracting translation patterns by comparing correct and wrong translations, as a means of enhancing a database of translation patterns. Arul Menezes and Stephen Richardson (“A Best-First Alignment Algorithm for Automatic Extraction of Transfer Mappings from Bilingual Corpora”) describe a logical-form alignment algorithm for the Microsoft MSR-MT architecture. Andy Way, in the final chapter (“Translating with Examples: The LFG-DOT Models of Translation”) describes various models based on adapting data-oriented parsing (DOP, a memory-based parsing method) to translation using an LFG-parsed bilingual treebank; he shows how this approach solves the problem of boundary friction (retrieved translations that do not fit the syntactic context).

Although it is very clear that the book consists of a number of independent papers and keeps some of the repetitive and overlapping flavor of a workshop proceedings volume, the editors and authors have done an excellent job in making this a coherent and self-contained book by adding cross-references and inviting additional papers. My main, not very vital regret regarding the book is that although there is some reference to relevant “lazy learning” techniques such as case-based, memory-based, and instance-based learning in AI, few links are made to the application of these ideas in areas of computational linguistics other than MT (ranging from case-based or memory-based phonology to pragmatics). As an example, see the special issue of *Journal of Experimental and Theoretical Artificial Intelligence* that I edited several years ago (Daelemans 1999). This work starts from the same inspiration as EBMT: Language-processing behavior is best modeled as similarity-based reasoning on the basis of stored experiences rather than as based on explicit rules extracted from these experiences. The motivation for this assumption, which has considerable empirical support, is that language data contain so many subregularities and exceptions that rules abstracting away from these exceptional or infrequent cases are at a disadvantage. It is my (undoubtedly biased) impression that this work sometimes goes a lot further than EBMT work in analyzing deeper questions such as why these techniques are better suited for NLP tasks and which task-independent similarity metrics and algorithms are best suited for solving NLP tasks in this paradigm. There may be a chance for mutually beneficial interaction here. In summary, I would recommend this book to everyone active or interested in MT, and especially the papers of the foundational part I to computational linguistics researchers in general.

References

- Aha, David W., editor. 1997. Lazy learning. Special quintuple issue, *Artificial Intelligence Review*, 11(1–5):1–423.
- Daelemans, Walter, editor. 1999. Memory-based language processing. Special issue, *Journal of Experimental and Theoretical Artificial Intelligence*, 11(3):287–471.
- Kasif, Simon, Steven Salzberg, David Waltz, John Rachlin, and David Aha. 1998. A probabilistic framework for memory-based reasoning. *Artificial Intelligence*, 104(1–2):287–311.
- Nagao, Makoto A. 1984. A framework of a mechanical translation between Japanese and English by analogy principle. In Alick Elithorn and Ranan Banerji, editors, *Artificial and Human Intelligence*, pages 173–180. Amsterdam: North-Holland.

Walter Daelemans is a professor of Computational Linguistics at the University of Antwerp (and part-time at Tilburg University). His research area is machine learning of language. Daelemans can be reached at CNTS, Department of Linguistics, Universiteitsplein 1 (A), B-2610 Antwerpen, Belgium; e-mail: walter.daelemans@ua.ac.be.

Current and New Directions in Discourse and Dialogue

Jan van Kuppevelt and Ronnie W. Smith (editors)
(University of Stuttgart and East Carolina University)

Dordrecht: Kluwer Academic
Publishers (Text, speech, and language
technology series, volume 22, edited by
Nancy Ide and Jean Véronis), 2003,
xii+381 pp; hardbound, ISBN
1-4020-1614-X, \$167.00, €152.00, £105.00;
paperbound, ISBN 1-4020-1615-8,
\$65.00, €59.00, £41.00

Reviewed by
Matthew Stone
Rutgers University

Van Kuppevelt and Smith's book offers a kind of archival proceedings for the Second SIGdial Workshop, which was held in 2001 in conjunction with Eurospeech. SIGdial is the Special Interest Group on Discourse and Dialogue of the Association for Computational Linguistics, and its annual workshop series, always collocated with major events in computational linguistics, has become a premier forum for empirical, formal, and computational approaches to language use. I think the breadth and energy of van Kuppevelt and Smith's collection supports an unequivocal appraisal of these meetings: If your interests touch on discourse and dialogue and your schedule and budget permit, you should be attending them. The next one will be held once again in conjunction with Eurospeech in September 2005 in Lisbon, Portugal. Go.

So should you buy this book? You might be worried about value. Inevitably, Kluwer wants \$167 from your library for the hardcover version. For that you get 12 revised versions of papers whose original versions are, in fact, archived in the ACL anthology. But before you go spooling things off to the printer, Kluwer does have a paperback version, on offer for a relatively modest \$65. Most of the chapters offer substantial additions (typically an extra methodological section) and spell out annotation standards, implementation techniques, or analytical results more thoroughly than the workshop papers. The expanded versions go a long way toward making the individual contributions more convincing and easier to replicate. In addition, the collection includes four invited new chapters, covering dialogue annotation, dialogue pragmatics, dialogue semantics, and dialogue system implementation. With all this new content, I wouldn't have hesitated to order the book for myself.

You might also be worried about the relevance of papers from a workshop from three years ago. In fact, the volume is as selective as any conference proceedings; van Kuppevelt and Smith report that the 12 contributed papers were winnowed down from 57 submissions to the workshop. However, rather than trying to present a concise and self-contained nugget of completed research, many of the most interesting papers articulate and motivate ambitious long-term agendas for dialogue research. The term *directions* in the book's title is thus faithful to its contents. The collection retains interest because it retains the adventurous and open-ended feel of a successful workshop and emphasizes these forward-looking characterizations of central problems and methods in the field. You can still use this book as a worthwhile jumping-off point for framing grant proposals or designing a graduate seminar in discourse and dialogue.

Three (invited) chapters nicely illustrate this flavor in the collection. All three chapters start from successful accounts of the dynamics of collaborative planning (Lochbaum 1998; Chu-Carroll and Carberry 2000), and take up the question of how such dynamics might be mediated by people's and systems' utterances. Jonathan Ginzburg uses corpus analysis to argue that the public contribution of an utterance follows closely from its semantics and is sharply distinguished from the speaker's more general implicatures and motives. While interlocutors' motives are always available as a metalevel topic of conversation, Ginzburg shows, only utterance content is at play in the fine-grained dynamics of clarification and grounding that allows interlocutors to achieve mutual understanding. This compelling new argument suggests that collaborative dialogue systems must signal their intentions publicly, directly in the content of their utterances. This is a challenge to semantics, as we must now represent such content in the "logical form" of discourse, perhaps along the lines being developed by Asher and Lascarides (2003). At the same time, Ginzburg's argument helps us to appreciate the conversational work that dialogue systems can do off the record, where the strategies dialogue agents use, such as indirection and politeness, can foster relationships with and trust among users (Bickmore 2003).

In another chapter, David Traum and Staffan Larsson survey their information-state approach to dialogue management. In essence, this approach offers a general knowledge representation methodology for use in characterizing dialogue context. It invites the designer to identify the meaningful distinctions in dialogue state that participants need to keep track of and to describe declaratively how utterances change it. Traum and Larsson's chapter recapitulates their earlier overview of the approach and survey of information-state systems (Larsson and Traum 2000) while clarifying their contribution to the design of reusable and sharable dialogue components. Although the chapter remains important, it might have been more instructive with a longer explanation of how and why to use the approach, as the approach continues to prove itself across an increasingly broad range of communicative behavior (Nakano et al. 2003) and increasingly varied interactions among communicators (Traum et al. 2003).

Lastly, Nate Blaylock, James Allen, and George Ferguson work to refine the moves of existing collaborative problem-solving models to the level of precision required to describe individual utterances in natural conversation. They build on previous models in factoring the course of collaboration into primitive joint steps, such as agreeing to adopt an action into a plan. (Their steps are particularly flexible because they draw on a somewhat formidable ontology to systematize the ingredients of collaborative planning and collaborative action.) They go on to characterize utterances by linking them with abstract communicative moves that specify one agent's contribution to one of these joint steps, such as initiating or completing one. Perhaps the most substantial insight here is that collaborative dialogues in dynamic domains play out differently because they must interleave planning and action. This is an increasingly important feature of current system domains from search-and-rescue to home automation. There is no going back. However, to go forward, we still need to regiment these collaborative moves to derive principled representations for agents' mental states, mutual commitments, and meanings—offering ample space for future research.

Obviously, an important theme of the collection is the interrelationships among human behavior, theoretical explanations, and system design. Two other themes of the volume deserve mention here as enduring inspiration for ongoing research in discourse and dialogue. The first theme is the challenge and promise of data annotation. For example, in the remaining invited chapter, Niels Ole Bernsen, Laila Dybkjær, and Mykola Kolodnytsky outline the requirements for annotation tools aimed at capturing the complexity of face-to-face conversation. As they outline, an ideal platform

would accommodate a range of users and uses and would support the analysis of communicative behaviors, including all aspects of their structure and interpretation, across all available modalities. This general vision, well-articulated here in one specific formulation, continues to set the agenda for tools and resources in projects like TalkBank (MacWhinney et al. 2004). Similarly, Lynn Carlson, Daniel Marcu, and Mary Ellen Okurowski document their work to create a corpus of discourse annotated with hierarchical interpreted structure in the framework of rhetorical structure theory. A set of case studies (new since the workshop submission) draw on their corpus to explore and characterize various levels of discourse: the uses of individual connectives within sentences, the expression of parallels and contrasts across multiple sentences, and the organization of documents as a whole. Annotated discourse corpora continue to be developed (Mitsakaki et al. 2004), so these strategies for mining and visualizing them will surely continue to inform their use.

The other theme is rapid development of dialogue applications. Two chapters in particular, one by Hiyun Alshawi and Shona Douglas, the other by Manny Rayner, Johan Boye, Ian Lewin, and Genevieve Gorrell (both slight revisions of original workshop papers), frame the issues in crisply insightful ways. Useful dialogue applications link domain-independent words and constructions with domain-specific responses. Configuring a dialogue system therefore requires specifications that crosscut traditional units of system design, such as grammatical resources and domain-general and domain-specific interpretive procedures. Building these specifications remains a complex task requiring enormous expertise. Rayner and colleagues aim to simplify these specifications by describing a “plug and play” architecture. System components are described in terms of their domain functionality. At the same time, they are linked with predefined hierarchical collections of linguistic resources that encapsulate the interface for that functionality. With these tools, new components can be characterized by compact specifications that can be integrated seamlessly and elegantly with an existing system. Alshawi and Douglas, by contrast, aim to learn these specifications. They start from examples pairing utterances with their domain-specific interpretations. They reconstruct the interpretation for a new utterance by directly manipulating the interpretation of a similar utterance drawn from this sample set. It testifies to the excitement of dialogue research that the basic problem of linking language to a world model admits such divergent approaches. Indeed, these approaches represent extremes of a huge design space, with alternative delineations of tasks and resources and with different demarcations between what to specify by hand, what to capture directly from data, and what to learn from annotation. This space remains an inspiration for dialogue research, including some of my own, for example (Stone et al. 2004).

Van Kuppevelt and Smith have put together an inclusive, timely, and significant collection. The book has small flaws—notably a few glitches in the typesetting of tables and of references in running text—but I find that my most severe criticism of it is simply that it is a book. In many respects, the closest parallel to it would be an issue of *AI Magazine* or *Communications of the ACM*: a selective forum with a broad audience for exciting, programmatic but precise discussion of important trends in research. This collection made me wonder whether, with its proliferation of lively SIGs and far-flung meetings, ACL might deserve an analogous series.

References

- Asher, Nicholas and Alex Lascarides. 2003. *Logics of Conversation*. Cambridge University Press, Cambridge.
- Bickmore, Timothy. 2003. *Relational Agents: Effecting Change through Human-Computer Relationships*. Ph.D. thesis, Massachusetts Institute of Technology.

- Chu-Carroll, Jennifer and Sandra Carberry. 2000. Conflict resolution in collaborative planning dialogues. *International Journal of Human-Computer Studies*, 53(6):969–1015.
- Larsson, Staffan and David Traum. 2000. Information state and dialogue management in the TRINDI dialogue move engine toolkit. *Natural Language Engineering*, 6:323–340.
- Lochbaum, Karen E. 1998. A collaborative planning model of intentional structure. *Computational Linguistics*, 24(4):525–572.
- MacWhinney, Brian, Steven Bird, Christopher Cieri, and Craig Martell. 2004. TalkBank: Building an open unified multimodal database of communicative interaction. In *Proceedings of the 4th International Conference on Language Resources and Evaluation*, Lisbon.
- Miltsakaki, Eleni, Rashmi Prasad, Aravind Joshi, and Bonnie Webber. 2004. The Penn discourse TreeBank. In *Proceedings of the 4th International Conference on Language Resources and Evaluation*, Lisbon.
- Nakano, Yukiko I., Gabe Reinstein, Tom Stocky, and Justine Cassell. 2003. Towards a model of face-to-face grounding. In *Proceedings of the 41st Annual Meeting of the Association for Computational Linguistics*, Sapporo, Japan.
- Stone, Matthew, Doug DeCarlo, Insuk Oh, Christian Rodriguez, Adrian Stere, Alyssa Lees, and Chris Bregler. 2004. Speaking with hands: Creating animated conversational characters from recordings of human performance. *ACM Transactions on Graphics*, 23(3):506–513.
- Traum, David, Jeff Rickel, Jonathan Gratch, and Stacy Marsella. 2003. Negotiation over tasks in hybrid human-agent teams for simulation-based training. In *Proceedings of the Second International Joint Conference on Autonomous Agents and Multi-Agent Systems*, pages 441–448, Melbourne.

Matthew Stone is an assistant professor in the Computer Science Department and the Center for Cognitive Science at Rutgers, the State University of New Jersey. Stone's 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>.