

Book Reviews

D. Jones, *Analogical Natural Language Processing*. UCL Press, London, 1996. ISBN 1-85728-218-3, £35.00.

Analogical Natural Language Processing aims to challenge the current hegemony of the rule-based paradigm in NLP. Traditional NLP decomposes languages into atomic units, whereas example-based NLP centres around the re-use of language fragments. The book consists of six chapters: a short introduction, chapters on background material, analogical machine translation, stochastic and analogy-based NLP, some experiments in analogical cloning and a conclusion. We shall now look at the chapters more closely.

The first, introductory, chapter outlines the scope of the book, namely analogical approaches (including the use of parallel distributed processing systems) to NLP and the suitability of a functional grammar representation for the task of storing example texts. Analogy-based machine translation, in particular the matching of source and target languages, is used as an example throughout the book. The basic analogical machine translation architecture, involving matching, recombination and generation is also presented.

The second chapter provides the background material which is necessary for the understanding of the subsequent chapters. The basic ideas behind example-based machine translation (EBMT), analogical modelling, functional grammar, parallel distributed processing, the PHRAN natural language front end and rhetorical structure theory are briefly introduced before the process of EBMT is looked at more closely. The decomposition/transfer/composition model of EBMT is presented, and both Sato and Nagao's and Sumita's recombination metrics are discussed, along with their pros and cons. The chapter concludes by examining functional grammar and its application to machine translation.

Chapter three centres on analogical translation. It covers the representation of translation examples and the nature of source examples, as well as the basic process of analogical translating. This involves using a shadow representation in functional grammar, linked to the example language text and using a match between these examples to provide the translation. The need to consider not only sentence structure but also text structure is made clear and amplified by means of examples. This not only leads to translations which are functionally closer to the rhetorical goals of the source text, but can be used to restrict the search for examples to those whose rhetorical function is closest to that of the original text. The problem with the need for the hand-coding of such relations is raised before a detailed example of EBMT is presented. Various problems are discussed at the close of the chapter, such as the problem with optional and obligatory functions, e.g. not all languages assign optional and obligatory functions in the same way, and a possible solution, making use of a hierarchy, is mentioned.

Chapter four looks at stochastic and analogy-based language processing. The precedents for such techniques are mentioned, e.g. word and sentence alignment in bilingual corpora, tagging parts of speech and other aspects of automated corpus analysis. Analogical modelling (AM) is presented as a general technique which lays down a challenge to structural methods. The preparation of analogical models and their use is discussed in detail and the AM technique is compared to PDP systems. Both are seen to offer graceful degradation and to merit further consideration.

This further consideration is given in the fifth chapter, which relates a set of four experiments in analogical cloning and compares the results obtained to those of PDP (Parallel Distributed Processing) approaches. The experiments involve determining the distance between prepositions, cloning verbs and their arguments, selecting between competing meanings and predicting case roles and lexical categories where these are not known. Finally, having covered the basic techniques, an example of EBMT recombination is presented, and this provides weight for the case for using analogical modelling in the place of a traditional rule-based parser and generator.

In the final chapter, the book concludes with a summary of the main arguments presented, namely that systems based on non-rule-based formalisms are powerful, that analogy has been shown to be a successful technique in the example domain of machine translation (in particular, that in rendering the generation phase a triviality it offers a fresh and efficient approach to machine translation), and that the analogical strategy bears some similarities to PDP systems. The chapter closes by highlighting a number of areas for further research, e.g. knowledge representation for examples, the relationship between recombination and analogical modelling in MT, the use of AM in other areas of MT, the applicability of RST and the computational cost of the algorithms on parallel machines.

Analogical Natural Language Processing offers a clear exposition of an novel technique. Whilst most of the examples in the book are Japanese, are solely concerned with machine translation and relate to relatively small experiments, the book makes interesting reading. There is a slightly disjointed feel to the book in places where the background material is not perhaps as fully expounded as it might be, or where the content seems strangely divided. For example, there is little insight, for the uninitiated, into the workings of PDP systems; the mention of PDPs at the end of chapter 4, and in more detail in chapter 5, seems oddly separated, and some of the experiments seem to have been tested on very small sets of data, especially for neural nets, thus making straightforward comparisons and analyses more tentative. These comments aside, the book still provides a good introduction to analogical modelling, and provokes thought as to how the technique could be employed in NLP in general. The book should interest all researchers in NLP, as well as being accessible enough for postgraduate course students.

CHRIS BOWERMAN

University of Sunderland

e-mail: chris.bowerman@sunderland.ac.uk

Lawrence A. Bookman, *Trajectories Through Knowledge Space: A Dynamic Framework for Machine Comprehension*. Boston, Dordrecht, London: Kluwer, 1994. ISBN 0-7923-9487-9, £62.95/\$89.85. 271 pp.

The availability of large text corpora has made statistical methods in natural language processing a feasible alternative or complementary approach to the ones used commonly in artificial intelligence. Problems encountered with symbolic hand-coded knowledge representations are related both to quantity and quality of knowledge. At the quantitative level it has been apparent that full-scale natural language interpretation presupposes large knowledge bases. A project like Cyc (Lenat *et al.*, 1990) has attempted to solve the problem by gathering a vast amount of knowledge. Dozens of person years have been spent to code in the rules, scripts, etc. It seems, however, that the qualitative problems are not solved by adding more and more pieces of symbolic knowledge representations to a system. Handling graded phenomena, fine-grained issues of polysemy, change and context-dependency, especially at the levels of semantics and pragmatics, are examples of problematic issues. Among others, George Lakoff (1987) has thoughtfully handled these matters from a linguistic and cognitive point of view.

Artificial neural networks (ANNs), inspired by the findings in human neurobiology, have

been presented to be a promising alternative to methods based on symbol manipulation. The necessary quantities of knowledge may be acquired from corpora by adaptive processes that are statistical in their nature. The numerical representations used in ANNs match more closely to the previously mentioned naturally imprecise phenomena in the mapping between language and our world. On the other hand, these methods have been criticised to lack, for instance, a natural means to represent dynamic structures.

Bookman's work is an important contribution considering the development of novel solutions for the adaptive and fine-grained natural language interpretation. It is based on the idea that the best results are gained by combining symbolic and connectionist representations. In the first two chapters, Bookman first gives an overview on the background of the LeMICON system. Using his words, LeMICON is a structured connectionist model that makes use of both connectionist and symbolic techniques to construct plausible interpretations of text. Bookman provides useful comparisons of symbolic, statistical, and connectionist approaches.

The main thrust of the book is built upon the LeMICON system. The basic building blocks of LeMICON include semantic memory with interconnected relational and associational tiers, and a working memory. Bookman introduces the term 'analog semantic features' (ASFs) that are derived from microfeatures. The third chapter of the book is devoted to the memory architecture of LeMICON and to the ASFs.

The fourth chapter is rather many-sided, perhaps even confusing. The basic algorithm is presented in some detail, and the central notions are introduced. For instance, each concept in semantic memory is represented by a vector of ASFs of dimension 454. These 454 named features are based on a selection from the category structure of the Roget's thesaurus (classes like 'payment', 'debt', 'credit', 'poverty', 'wealth', 'investment', 'finance', to mention some). Bookman also provides links to related psychological and neurophysiological research. Here the presentation is rather shallow in some respects which is, on the other hand, understandable taking into account the overall breadth of the work. The author also interestingly compares LeMICON to several other text understanding systems from different paradigms. Among other things, he considers capabilities related to word-sense disambiguation, integration of syntax, semantics and pragmatics, and the use of soft constraints.

The subsequent chapters show examples of actual processing with analysis of the interpretation results. Chapters 7 and 8 describe experiments in acquiring knowledge from on-line corpora and their analysis. Bookman demonstrates how to perform automatic knowledge acquisition using a 7 million word database of Wall Street Journal articles. The semantic memory of LeMICON is constructed automatically. One important realistic aspect of the system is that it dynamically changes the semantic memory as new texts are given as input.

The analysis of five factors related to the knowledge representation and acquisition of knowledge brings two recent connectionist systems to the top, namely DISCERN by Risto Miikkulainen (1993) and LeMICON. One main contribution of Miikkulainen is the application of Kohonen's self-organising maps (Kohonen, 1995) in lexical processing to enable unsupervised learning while processing the input examples. Also Bookman emphasises unsupervised processing that provides one cornerstone for the scalability of the system.

At the end of the book, the author discusses the limitations of his approach and possible future directions. According to Bookman, the main problems are related to the inability to deal with metaphoric text, the lack of knowledge of how social role constraints affect interpretation, the limited form of temporal reasoning, the lack of integration of syntax and semantics, and finally, the fact that LeMICON must preprocess the background knowledge for each new text.

This work deserves serious attention from the natural language processing research community. The result is a good compromise and mediator between more radical approaches of polished artificial neural networks, and traditional symbolic AI methods. Bookman's work may also provide many insights for those who are developing practical systems for NLP applications. In the long run, one may foresee development of more radical theories and systems that include, for example, explicit visual information, and other forms of inherently

non-symbolic modalities. Bookman's work may be seen to be a step towards a position where the complex cognitive and social phenomena related to full-scale natural language processing are taken into account. The development of successful systems that possess near-to-human communication ability may require thoughtful consideration of Heinz von Foerster's (1981; originally 1972) formulations: a formalism necessary and sufficient for a theory of communication must not contain primary symbols representing communicabilia. The nervous activity of one organism cannot be shared by another organism. Communication is an observer's interpretation of the interaction between two organisms.

References

- von Foerster, H. (1981) *Observing Systems*. Seaside, CA: Intersystems Publications.
 Kohonen, T. (1995) *Self-Organizing Maps*. Berlin: Springer-Verlag.
 Lakoff, G. (1987) *Women, Fire, and Dangerous Things: What Categories Reveal About the Mind*. Chicago: University of Chicago Press.
 Lenat, D. B., Guha, R. V., Pittman, K., Pratt, D. and Shepherd, M. (1990) Cyc: Toward programs with common sense. *Communications of the ACM* **33**: 30–49.
 Mäikkyläinen, R. (1993) *Subsymbolic Natural Language Processing: An Integrated Model of Scripts, Lexicon, and Memory*. Cambridge, MA: MIT Press.

TIMO HONKELA

Helsinki University of Technology

e-mail: Timo.Honkela@hut.fi

John Lyons, *Linguistic Semantics: An Introduction*. Cambridge, UK: Cambridge University Press, 1995. ISBN 0 521 43877 2 (paperback). ISBN 0 521 43302 9, Price £35.00 (hardback). xvii + 376 pages.

Sir John Lyons's *Linguistic Semantics: An Introduction* (LSAI from now on) is a tolerable addition to the list of half a dozen or so impressive titles he has produced on linguistic subjects over the years. This book was initially planned to be a second edition of his *Language, Meaning and Context* (Lyons, 1981). However, in the end it turned out to be a successor and replacement. For it is, in the author's words, a very different book compared to the 1981 volume: it is much longer, treats topics missing in the earlier volume, and is written in a different style. (Unfortunately, I am not familiar with Lyons (1981), and the reader is asked to take these remarks with a grain of salt.)

By 'linguistic semantics' Lyons means the study of meaning systematically encoded in the vocabulary/grammar of natural language. Thus, linguistic semantics is a branch of linguistics; semantic issues which have more to do with philosophy belong, in Lyons's view, to the more proper branch of philosophical semantics. Accordingly and understandably, Lyons devotes limited space to philosophical problems while he cautions that nobody would be able to appreciate modern linguistic semantics without some acquaintance with its philosophical groundwork.

According to Lyons, LSAI can be used as a textbook for introductory semantics courses in departments of linguistics. If one is familiar with various key issues in semantics, then this book is quite enjoyable because Lyons is raising some interesting points and asking stimulating questions. On the other hand, a beginner would be easily unsettled by the lack of exercises, light (at times skin-deep) treatment of some very significant topics of semantics, and the generally verbose writing style of Lyons. In fact, this last point needs some emphasis. To repeat the words of Wittgenstein *vis-à-vis* a well-known philosopher (Rhees, 1984, p. 88): "He is too long-winded; he keeps on saying the same thing over and over again. When I read him I always wanted to say, Oh all right, I agree, I agree, but please get on with it."

It should be added that since LSAI is not a *formal* semantics book, it is natural that the writing style suffers from a good deal of repetition – probably thought of as a cure for ambiguity. In a formally written book, the mathematics would take care of the rest, and less prose would be needed. However, in LSAI there is only a minimal amount of mathematics. (This, by the way, reminds me of *Daedalus*, an influential American journal which allows no formulas.)

LSAI consists of the following parts (each part is followed by the chapters comprising it):

1. Setting the scene
 - Metalinguistic preliminaries
2. Lexical meaning
 - Words as meaningful units
 - Defining the meaning of words
 - The structural approach
3. Sentence-meaning
 - Meaningful and meaningless sentences
 - Sentence-meaning and propositional content
 - The formalization of sentence-meaning
4. Utterance-meaning
 - Speech acts and illocutionary force
 - Text and discourse; context and co-text
 - The subjectivity of utterance

Part 1 (which consists of one long chapter) deals with the Herculean question of semantics, linguistic or non-linguistic: what is meaning? The need for a metalanguage of semantics is elucidated and Standard English is adopted as the metalanguage. This chapter also introduces the usual distinctions between language and speech, *langue* and *parole*, competence and performance, form and meaning, sentences and utterances.

The first chapter of Part 2 is concerned primarily with words as expressions. The fundamental distinction of C. S. Peirce, viz. tokens vs. types, is explained. Homonymy, polysemy and synonymy are studied in long expositions. There is also an interesting discussion of full (e.g. *man*, *green*) vs. empty (e.g. *the*, *to*) word-forms. The second chapter of Part 2 enumerates the different techniques that can be employed to define the meaning of words. Classical notions such as sense, reference, extension, intension, ostension, natural kinds and prototypes are introduced. The final chapter of Part 2 hints at the notion of compositionality. It also studies (propositional) entailment while drawing the classical distinctions between necessarily vs. contingently true propositions. As a result of reading this part, it would be nice if the reader got an overall picture of what word meaning is all about. Unfortunately, this does not happen. Personally, I might as well prefer the following holistic explanation (Suzuki, 1984, p. 79):

“Consider the word *dog*. The word *dog* evokes very different images in the minds of those who have been bitten by dogs and those who have not. Dog lovers can distinguish different kinds of dogs and point out their respective characteristics, while people who are not interested in dogs cannot. But this does not keep anyone from understanding and using the word *dog*. With this in mind, I feel that the meaning of a word should be as follows: the sum total of all the individual experience and knowledge we have in connection with a certain phonological shape.”

While Part 2 has been concerned with lexical semantics, Part 3 moves on to a consideration of sentence-meaning. The verifiability principle of A. J. Ayer (“A sentence is factually significant to a given person if and only if he knows how to verify the proposition which it purports to express”) is given and emotivism (the thesis that statements in ethics and esthetics just codify one’s feelings and are not saying anything that is true or false) is presented. The second chapter of this part has a fine discussion of implication and negation in natural language utterances.

The meaning of interrogative as well as imperative and exclamative sentences are studied, in addition to the usual declaratives. The last chapter of Part 3 details the principle of compositionality. It also presents a dated theory (which Lyons dubs ‘Katz-Fodor theory of sentence meaning’) together with the time-honored distinction about deep vs. surface structure, and selection-restrictions for handling semantic ill-formedness. There is also a gainless discussion of Montague grammar and possible worlds which would probably cause more confusion for a beginner than clarification. It is a pity that Montague grammar, this particularly fruitful approach to the analysis of natural language, is subjected to such a light treatment in the hands of Lyons. To be fair, Lyons acknowledges that Montague grammar is a very technical subject and that he wishes to explain, non-technically, only some of the most essential aspects of it. (In this endeavor he succeeds to a certain extent.) My overall feeling is that Lyons contends that the standard view of logic is inappropriate for natural language semantics. While I largely concur with this message, I still think that the following excerpt from Strawson (1985, p. 232) does better justice to logic:

“[In] trying to discover the answers to questions of such forms as ‘What are the conditions under which we use such-and-such an expression or class of expressions?’ or ‘Why do we say such-and-such a thing and not such-and-such another?’, we may find ourselves able to frame classifications or disclose differences broad and deep enough to satisfy the strongest appetite for generality. What we shall not find in our results is that character of elegance and system which belongs to the constructions of formal logic. It is none the less true that the logic of ordinary speech provides a field of intellectual study unsurpassed in richness, complexity, and the power to absorb.”

Part 4, based on J. L. Austin’s theory of speech acts, includes some of the best chapters of the book. The chapter on contexts is particularly notable in this regard. When the word *context* is employed in linguistic descriptions, explanations and analyses, its meaning is frequently left to the reader’s understanding, i.e. it is used in an implicit and intuitive manner. In fact, the denotation of the word has become murkier as its uses have been extended in many directions; in short, context has become some sort of conceptual garbage receptacle. Lyons offers, using Paul Grice’s accounts of conversational implicature, principle of cooperation and conversational maxims, a profitable review of the role of context. His feeling is that in the construction of a satisfactory theory of context, the findings of social sciences like psychology, anthropology and sociology will be useful. Finally, the last chapter of this part considers deixis, tense and aspect, and mood. It is a difficult chapter that emphasizes what Lyons calls the ‘subjectivity of utterance’.

Now that I have given a taste of what is in this book, it is time to voice my disappointment with various shortcomings and omissions.

Government and Binding Theory and *the Minimalist Program* (Chomsky, 1995) receive scant attention in LSAI. Lyons’s most recent Chomsky citation is from 1986. I find this hard to understand in the light of the fact that Lyons is also the author of a monograph on Chomsky (Lyons, 1991).

Situation Semantics (Barwise and Perry, 1983) and *Discourse Representation Theory* (Kamp and Reyle, 1993) are not even mentioned. I find this altogether strange. Obviously, both theories have been around not to be regarded as exotic research fads, and have given rise to landmark publications and practical applications in semantics and natural language processing, respectively.

Donald Davidson, whose contributions to the truth-conditional theory are numerous, is another researcher who is simply ignored in LSAI (not counting a passing reference made in ‘Suggestions for further reading’). As is well known, Davidson holds that Tarski-style truth theories are fundamental: his standpoint is that there is a close affinity between understanding the meaning of a sentence and knowing its truth conditions. On a related note, while Kaplan’s work (Kaplan, 1989) on the logic of demonstratives has had a profound reign in semantics, it is simply overlooked in LSAI.

Still, Lyons should be commended for writing this book. Writing books is risky business; one does one's best and lets the result stand or fall. In my view, LSAI does not fall but leaves a lot to be desired. As long as one can bear with the sometimes odd (even eccentric) points of view of Lyons, there are no obvious 'conceptos'. But there are a small number of typos, e.g. in the Venn diagram on page 111 $\sim A$ should read A , and on page 167 the formula $(\sim p \ \& \ \sim p)$ should read $(\sim p \ \& \ \sim q)$.

In my own graduate teaching (a course titled *Varieties of Formal Semantics*), I used to prefer Chierchia and McConnell-Ginet (1990), keeping Gamut (1991) as a supplementary source to provide the necessary background material as needed. Both of these books are more or less formal and include exercises. I think I'll stick to my choices, but will probably ask my students to take a look at LSAI. After all, inverting a remark of Alice ter Meulen, "This is not the textbook on semantic theory that the world has been waiting for!"

References

- Barwise, J. and Perry, J. (1983) *Situations and Attitudes*. Cambridge, MA: MIT Press.
- Chierchia, G. and McConnell-Ginet, S. (1990) *Meaning and Grammar: An Introduction to Semantics*. Cambridge, MA: MIT Press.
- Chomsky, N. (1995) *The Minimalist Program*. Cambridge, MA: MIT Press.
- Gamut, L. T. F. (1991) *Logic, Language, and Meaning* (Volume 1: *Introduction to Logic*, Volume 2: *Intensional Logic and Logical Grammar*). Chicago: University of Chicago Press. (First published 1982 (in Dutch).)
- Kamp, H. and Reyle, U. (1993) *From Discourse to Logic: Introduction to Model-Theoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory*. Dordrecht, Netherlands: Kluwer Academic.
- Kaplan, D. (1989) Demonstratives. In J. Almog, J. Perry and H. Wettstein (eds.), *Themes from Kaplan*. New York: Oxford University Press. (Article existed in manuscript form since 1977.)
- Lyons, J. (1981) *Language, Meaning and Context*. Glasgow, UK: Fontana/Collins.
- Lyons, J. (1991) *Chomsky*. London: HarperCollins. (Third, revised and further enlarged edition. First published 1970.)
- Rhees, R. (ed.) (1984) *Recollections of Wittgenstein*. Oxford, UK: Oxford University Press.
- Strawson, P. F. (1985) *Introduction to Logical Theory*. London: Methuen. (First published 1952.)
- Suzuki, T. (1984) *Words in Context: A Japanese Perspective in Language and Culture*. New York: Kodansha International. (Translated by A. Miura. First published 1973 (in Japanese).)

VAROL AKMAN
Bilkent University, Ankara

e-mail: akman@cs.bilkent.edu.tr