

Briefly Noted

Language and Learning for Robots

Colleen Crangle and Patrick Suppes
(Stanford University)

Stanford, CA: Center for the Study of Learning and Information (CSLI lecture notes 41), 1994, xxi+276 pp; distributed by the University of Chicago Press; hardbound, ISBN 1-881526-20-8, \$49.95, £39.95; paperbound, ISBN 1-881526-19-4, \$21.95, £17.50

In a career spanning five decades, Patrick Suppes has pursued a unique approach to the problems of language and learning. This book, with Colleen Crangle, is a reprise of much of the work on the philosophy of language. As the title indicates, the underlying model is of a robot that is to receive instructions (and later also instruction) in ordinary language. The philosophizing is made concrete through software simulations and some simple experiments with a physical robot. Although much of the material has appeared before, it is reworked into a cohesive whole.

Any attempt to deal with such deep issues will involve simplifications, and it is instructive to follow the unfolding of complications. The starting point is quite conventional; attribute grammar with a semantics based on formal logic. But the demands of the robot domain quickly require the introduction of functional semantics for perception and action, which are not analyzed further, and for control primitives. Task requirements also push the authors to abandon the autonomy of syntax and to give a major role to intentionality, contextual readings, and the shared perceptions of speaker and hearer. Thus they are led to many of the core concerns of cognitive semantics and this becomes explicit in a chapter on geometry and spatial prepositions. Traditional issues concerning negation get extended in intriguing ways in the chapter that considers what it means to ask a robot to stop.

The work on learning is mostly recent, although the theme goes back to Suppes's original work in mathematical psychology. The model has the robot choose parameters of actions (distance, etc.) from probability distributions that get modified by verbal feedback. The final chapter describes on-going work on the learning of natural language from examples in the robot domain. By postulating

pre-existing internal structures for each utterance, the problem is reduced to statistically associating surface forms with internal constructs. This problem is analyzed theoretically and on a small corpus. As in the rest of the book, the work is idiosyncratic. But anyone with a deep interest in the philosophy of language can benefit from following the progression of ideas on some of its most fundamental issues.—*Jerome A. Feldman, International Computer Science Institute, Berkeley*

Semantics: Defining the Discipline

Robert A. Hipkiss
(California State University, Long Beach)

Mahwah, NJ: Lawrence Erlbaum Associates, 1995, xvii+123 pp; hardbound, ISBN 0-8058-2026-4, \$29.95; paperbound, ISBN 0-8058-1593-7, \$17.50

This short book is meant to introduce the discipline of semantics, which Hipkiss believes ought to be recognized as *sui generis* and not a component of linguistics, philosophy, psychology, or anything else. Nevertheless, he reveals a linguistic bias by defining semantics as "the study of word meanings" (p. 107). To create an insightful, usable introductory semantics text of a mere 130 pages (introduction and 10 chapters) would be a masterstroke. It is very difficult to write an introductory text of any length for a wide-ranging topic like semantics (in contrast with set theory or predicate calculus). The book begins with gusto, and moves from Plato to Wittgenstein within three pages, touches some aspects of linguistic semantics, and ends with a paean to Korzybski's general semantics—an ideological position on meaning in language that current theoretical semantics largely (and properly) ignores. There are too many significant gaps and factual errors for this book to be acceptable; I felt I was reading an essay by an enthusiastic undergraduate with many flashes of common sense but too little learning.

Some examples of naive and misleading or incorrect statements, and failure to discuss important matters arising will ground these criticisms:

- Chapter 8 asks "Is there a semantics for AI?" and answers "AI is showing that the computer is far more than a mechanical

memory. It can learn and conceptualize. It can relate those concepts to language. It can, then, give semantic content to the words in its memory. This is a limited semantics, however, for it has, at best, a limited intentionality" (p. 96). I wish my computer could conceptualize!

- "[F]eature analysis selection [is] the favorite subject of Fillmore and other case-grammarians" (p. xv). (!?)

- "Kripke (1980) sees the Morning and Evening Star as the same, contra Russell and Frege, but that is because he looks at the object by itself as though continually observed, not at it in its contexts in different manifestations" (p. 30).

- Hipkiss espouses Wittgenstein's meaning-as-use theory without discussing its ramifications, and implies that a semanticist requires no special training: "Ultimately we are forced to do what the lexicographers do, listen, read, and learn how words are actually used, in order to know what they mean to different people. The more we read and listen and practice different modes of discourse at various levels of abstraction the more finely tuned we become to the possible meanings of different words in different contexts. There is no substitute for that experience, although our dictionaries, to some extent, abridge that experience for us" (p. 107).

- Common sense leads Hipkiss astray: "Normally, we would not run through all the dictionary meanings of verb and noun when we see a word, for we have already been limited in our search by the context of the sentence" (p. 69). In fact, all senses *are* activated and then inappropriate ones suppressed—see Gernsbacher (1990) and references cited there.

- The claim that "propositions are, by definition, true statements" (p. xiii) needs discussion in an introductory textbook; at worst it makes a mockery of propositional logic, and at best should lead to discussion of what is meant by a true versus a false statement.

- It is incorrect that around 1970 Jerrold Katz "developed in his generative semantics a system of semantic marking that explained, to some degree, the mental process of selecting various words and word order as we speak and read" (p. xiv). Katz was never a generative semanticist and although he once believed he was defining semantic competence in Chomsky's sense, and later the se-

mantics of an abstract system, he specifically eschewed performance and processing models (see Katz 1966, 1972, 1981).

- Erroneous is: "Wierzbicka (1991) created a pragmatics dictionary of English verbs" (p. 39). The correct reference is Wierzbicka 1987, and it presents the senses of speech act verbs, rather than their pragmatics.

- Chapters 1–3 discuss conceptualization, subjective awareness, and relative truth, while dismissing "logic" as a means of accounting for these important aspects of semantics. Although Hipkiss writes "words . . . are symbols for concepts arrived at through experience" (p. 2), there is no discussion of cognitive or conceptual semantics.

- An insufficient grasp of logic seems to lie behind a number of statements in this book—for example, the discussion of quantification and the definite article (p. 16), and "were truth-conditional semantics to incorporate the possible and the probable, it would not have been such a narrow incursion into the worlds of language and thought" (p. 26). Hipkiss claims that possible-worlds semantics is needed, but contrasts it with truth-conditional semantics, thus ignoring its truth-conditional basis. Nowhere is model-theoretic semantics so much as mentioned.

- There is far too little discussion or explanation in this book. For example, cases (thematic roles) are listed without exposition (p. 63). Chapter 9 lists titles of papers in *Etc.*, the flagship journal of general semantics, but there is almost no discussion of them, nor any comparison of general semantics with theories of semantic structure and relations.

Such infelicities are just a sample. This is a book that should be approached with caution by anyone with a limited (or no) knowledge of semantics. True, no book is perfect, but too much of this one is either misleading or just plain wrong.—Keith Allan, Monash University

References

- Gernsbacher, Morton Ann. 1990. *Language Comprehension as Structure Building*. Hillsdale, NJ: Lawrence Erlbaum.
- Katz, Jerrold J. 1966. *The Philosophy of Language*. New York: Harper and Row.
- Katz, Jerrold J. 1972. *Semantic Theory*. New York: Harper and Row.
- Katz, Jerrold J. 1981. *Language and Other Abstract Objects*. Totowa: Rowman and Littlefield.

- Kripke, Saul. 1980. *Naming and Necessity*. Oxford: Blackwell.
- Wierzbicka, Anna. 1987. *English Speech Act Verbs: A Semantic Dictionary*. Sydney: Academic Press.
- Wierzbicka, Anna. 1991. *Cross-cultural Pragmatics: The Semantics of Human Interaction*. Berlin: Mouton de Gruyter.

Reasoning about Knowledge

Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi
(IBM Almaden Research Center, IBM Almaden Research Center, Weizmann Institute of Science, and Rice University)

Cambridge, MA: The MIT Press, 1995,
xiii+477 pp; hardbound, ISBN 0-262-06162-7,
\$45.00

Reasoning about knowledge is concerned with answering questions like: "What do I / you / we know?", "What do I / you / we not know?", and "Do I / you / we know enough to..." Reasoning about knowledge underlies all of human communication, as communication is an attempt to make someone else know (or believe) something. Further, communication depends on what the participants know about each other's knowledge. Reasoning about knowledge is also of interest in cryptography, theory of economics, game theory, philosophy, and general artificial intelligence.

The book *Reasoning about Knowledge* analyzes reasoning about knowledge held by groups of agents. (The book uses this currently-overused term, so I will too.) The authors are among the foremost experts in the area and are responsible for many of the recent advances in reasoning about knowledge.

Reasoning about Knowledge emphasizes perfect reasoning, which is not appropriate for actual human linguistic systems, but an understanding of the simpler, perfect reasoning case is needed, in my opinion, as a precursor to the more complex, limited or imperfect reasoning case, and the book does devote considerable space to some *formal* models of the limited or imperfect reasoning case. (Whether formal treatments of the more complex case are appropriate is a separate matter; even if you think such formal methods are garbage, it is generally a good idea to know

something about rival camps.)

The book presents the formal underpinnings of reasoning about knowledge. It covers the defining of several varieties of multi-agent systems, which are systems composed of several agents that perform actions and use reasoning about knowledge or, perhaps, just act like they do or can profitably be analyzed as such. The book discusses common knowledge and its relationship to agreement, including the paradox that common knowledge is needed for agreement but is essentially impossible to achieve. The book mentions some of the problems with logical omniscience and presents some formal systems that are not logically omniscient.

There are many examples in the book that make its content much easier to understand. The book is primarily written for use in courses, but would be useful for anyone with an interest in reasoning about knowledge and with some background in formal logic. It contains many exercises suitable for classroom use or self-study.

In summary, this is an extremely well-written book that covers a subject that is important to computational linguistics. I recommend it for anyone interested in situations that occur when one agent must reason about the knowledge of another.—*Peter F. Patel-Schneider, AT&T Research*

Empirical Methods for Artificial Intelligence

Paul R. Cohen
(University of Massachusetts)

Cambridge, MA: The MIT Press, 1995,
xvi+405 pp; hardbound, ISBN 0-262-03225-2,
\$55.00

Paul Cohen has written a remarkable and valuable guidebook for those doing empirical research in artificial intelligence and other computational fields. This book differs in flavor from those devoted to experimental design or statistical methods in that all techniques are set firmly within motivated, compelling discussions of the scientific method and how it can be properly applied when the object of study is a complex computer program.

The book is organized as a progression through the experimental process, and

brings together, in clear, approachable prose, discussions of research methodology, exploratory data analysis, experimental design, hypothesis-testing tools (traditional and computer-intensive), causal modeling of system behavior, and techniques for generalizing from data and analyses. Examples and case studies are abundant, drawing most heavily from real experiments in computational linguistics, machine learning, and planning.

Throughout the book, strategies for how to design and understand the experimental tasks are foremost; statistics are treated as tools for achieving these ends. For example, topics such as analysis of variance and factorial design, which are chapter headings in many texts, are discussed here in the context of how to explain system performance.

Computational linguists should be aware that despite its title, this book contains little discussion of probability, multidimensional scaling, learning algorithms, and other topics that are important for corpus-based language analysis but not necessarily for experimental design. Additionally, readers not familiar with AI should be warned that the book contains some AI-centric examples whose motivation is not well-explained, as well as frequent references to debates within AI on how to conduct and evaluate experiments. Another potential drawback, if one wishes to use this book as an instructional text, is the lack of practice problems. However, one could exercise the concepts learned by having students attempt to apply them to published research papers. The book begins at an introductory level, assuming no background in statistics, and although it develops many advanced techniques, once this material is mastered, supplementary references may be required (and Cohen usually supplies relevant pointers).

One could argue that there is no need for such a book, given the legions of existing expositions on scientific methodology and experimental design. However, *Empirical Methods for Artificial Intelligence* is useful not only because it is targeted at a computational audience, whose training so frequently and unfortunately does not include basic research methodology, but also because it presents the material in a clear, useful, modern, and inviting manner.—*Marti A. Hearst, Xerox Palo Alto Research Center*

Applied Logic: How, What and Why. Logical Approaches to Natural Language

László Pólos and Michael Masuch (editors)
(Centre for Computer Science in Organization and Management, University of Amsterdam)

Dordrecht: Kluwer Academic Publishers
(Synthese Library, edited by Jaakko Hintikka, volume 247), 1995, viii+392 pp; hardbound, ISBN 0-7923-3432-9, \$115.00, £74.00, Dfl 175.00

This book arises from the first Conference¹ on Applied Logic organized in December 1992 in Amsterdam by the Centre for Computer Science in Organization and Management and the Institute for Logic, Language, and Information. The collection contains 13 papers that were classified by the program committee as belonging to the area of "Applications of logic to the study of natural language."² The editorial preface is extremely brief. No overview, comparative summary, or evaluation of the contribution of the individual papers is provided.

I would call the papers thin on data and heavy on theory. Only a few briefly address computability. The authors appear to be much more concerned and familiar with logic than with natural language. Two papers, "Optimization of deduction for multimodal logics" by Olivier Gasquet and "Implicit and explicit definability in modal and temporal logics" by Larisa Maksimova, do not discuss any data and do not contain any discussion about natural language whatsoever. Only one paper, "Dynamic aspect trees" by Jerry Seligman and Alice ter Meulen, discusses some real-looking data (although its source is not indicated), including a multi-paragraph text.

Reading and understanding such extremely formal, long papers (all but one between 20 and 40 pages), is not an easy task, even for one accustomed to formal theories. The papers are journal-like—many report on mature research—but they do not appear to

1 The editors also use the term *workshop* when referring to this meeting.

2 The other two topic areas were "Computer science applications of logic" and "Algebraic applications."

have undergone a rigorous journal-like review.

The reader is often left wondering whether the ideas presented are new or whether they were reported elsewhere. Sometimes simple-to-fix stylistic problems, such as claims buried toward the end of the paper, significantly obscure the presentation. The omission of some very relevant references is often apparent, with few references to the work outside Europe.³ For example, the paper "Are types needed for natural language?" by Fairouz Kamareddine, surprisingly does not reference Keenan and Faltz (1985). Many papers discuss dynamic characteristics of their formalisms. The word *dynamic* here refers to step-by-step text interpretation accompanied by appropriate updates of information obtained from the text seen thus far. It would be interesting to know how these dynamic logics differ from some existing formal frameworks for processing natural language that view understanding of natural language as the process of knowledge update. It would also be interesting to know how these dynamic logics differ from the knowledge representation formalisms, including belief-revision systems, that have been proposed over the last decade in artificial intelligence.

While I liked fragments of many papers, I was generally not convinced by the arguments in support of the claimed advantages of the formalisms and even less so about their relevance or criticality for processing natural language.

My recommendations are as follows. I would not recommend this book to the general computational linguistics audience. Researchers who are concerned with handling large volumes of real linguistic data are highly unlikely to benefit from theories motivated by such artificial, toy examples as those discussed in this book. More-formally inclined researchers will probably find only selected papers to be of interest. I would recommend this book to those unfamiliar with European research on applied logic; 14 out of 18 authors are affiliated with European institutions and their research provides a good sample. I would also recommend this book

as a source of logical formalisms with some relevance to natural language. While one's chances of learning about results of immediate interest are small, there is a great opportunity to extend one's horizon. For example, this book has greatly improved my limited knowledge of Situation Theory.⁴ Similarly, researchers not familiar with the theory of programming languages can learn some common themes in the semantics of programming and natural languages.—*Lucja Iwańska, Wayne State University*

Reference

- Keenan, Edward L. and Leonard M. Faltz.
1985. *Boolean Semantics for Natural Language*.
Dordrecht: D. Reidel.

HORATIO: A Middle-sized NLP Application in Prolog

Archibal Michiels
(Université de Liège)

Liège: Department of English, Université de Liège, 198 pp and diskette; paperbound, ISBN 2-87233-015-i, BEF 700
[Available from: L3, Dept of English, Université de Liège, place Cockerill 3, B-4000, Liège, Belgium, fax +32-41-665721; add BEF 50 for postage outside Belgium.]

HORATIO is a moderately-wide-coverage English parser developed in Arity Prolog on a PC and based on Michael McCord's slot grammars. Its output includes explicit grammatical relations such as subject and object, and some lexical disambiguation is performed by using argument properties such as *human* and *abstract*. Performance is respectable; on my 66-MHz 486, each sentence in the test suite took no more than a few seconds, including complex sentences such as *Devito manages a programmer Abrams interviewed and Browne hired*.

Noteworthy features of HORATIO include a lexicon derived (by the *awk* program) from an on-line version of the *Longman Dictionary of Contemporary English*; a parsing technique

³ It is only fair to note that many American researchers are similarly guilty of largely ignoring European results.

⁴ Which is not to say that I like this theory. It appears to offer very little for representing meaning of natural language sentences involving general negation.

that avoids looping on coordinate structures; provisions for idioms with different degrees of syntactic opacity; a mechanism to handle word-order variation; and a generation facility.

The book documents HORATIO in detail: "No pseudo-code is given. No piece of code is presented in a simplified form. Cited code always corresponds exactly to the runnable code to be found on the companion disk" (p. 5).¹

The diskette does indeed contain the source code, executables, and test suite, entire. The user can verify that HORATIO works as advertised, give it additional sentences to parse, and use it as the basis of further work (taking due note of the copyright, of course). Porting from Arity Prolog to other implementations should be straightforward.—*Michael A. Covington, University of Georgia*

Language Industries Atlas

Paul M. Hearn and Diana F. Button
(editors)
(INK–Belgium)

Amsterdam: IOS Press, 1994, xviii+406 pp; hardbound, ISBN 90-5199-148-7, no price listed

"The aim of our study has been to describe the activities of the many organisations, both public and private, that create the infrastructure within which languages are able to develop and interact on equal terms in multilingual Europe. Although developments in the United States and Japan have been mentioned, we have been interested primarily in describing activities in the European Community. . . .

"The *Language Industries Atlas* will be of interest to a broad spectrum of language professionals including translators, interpreters, terminologists, teachers, computational linguists, researchers and public and private organisations, and indeed to anyone interested in languages, from a planning, standards, infrastructural or technological perspective.

"The *Atlas* contains more than 1,000 descriptions of activities which play a role in

shaping the language industries, whether from a user or provider perspective. . . . [It] has been compiled from a questionnaire and telephone survey carried out over the past 12 months. . . . Where information was not provided from our survey, we have supplemented this from our own records."—*From the introduction*

Text Encoding Initiative: Background and Context

Nancy Ide and Jean Véronis (editors)
(Vassar College and CNRS and Université de Provence)

[Reprinted from *Computers and the Humanities*, volume 29, numbers 1, 2, and 3, 1995.]

Dordrecht: Kluwer Academic Publishers, 1995, vi+242 pp; hardbound, ISBN 0-7923-3689-5, \$99.00, £60.00, Dfl 145.00

"The Text Encoding Initiative *Guidelines for Electronic Text Encoding and Interchange*. . . , known familiarly as TEI P3 . . . , appeared in May 1994. The Guidelines are the result of over six years' work by dozens of scholars from all over the world who were involved in TEI working groups, providing their conclusions concerning the optimal way to consistently and comprehensively encode a vast range of text types and features. As such, TEI P3 represents a pioneer effort in an area where only occasional and isolated attempts had been made before, and will certainly serve as the primary basis for encoding texts in electronic form for the foreseeable future.

"There is a compelling reason for collecting this series of papers. The TEI is a pioneering effort; the TEI working groups were the first to comprehensively address the substantial intellectual problem of representing textual data in electronic form. They were faced with a vast array of new problems: it became immediately apparent that the development of a text encoding scheme demanded much more than assigning tag names to features, and included looking at the conceptual structure of texts and determining the commonalities across different text types. It also demanded finding the most consistent and effective ways to encode texts using the Standard Generalized Markup Language (SGML), which provides only the fundamental machinery for marking texts and provides

¹ The code may also be downloaded free of charge from <http://engdep1.philo.ulg.ac.be>.

for many alternative means to encode texts. Therefore, the work of participants in the TEI not only involved consideration of problems of text encoding that are likely to be with us for decades to come, but also required the development of a methodology, from scratch, for approaching these problems. These pioneering efforts, while likely to be refined and extended, should not be lost; they provide the intellectual basis upon which text encoding practices will build in the future. This collection is therefore also an attempt to document the course of these efforts."—*From the publisher's announcement*

Elements of Acoustic Phonetics (second edition)

Peter Ladefoged

(University of California, Los Angeles)

Chicago: The University of Chicago Press, 1996, viii+216 pp; hardbound, ISBN 0-226-46763-5, \$39.95, £31.95; paperbound, ISBN 0-226-46764-3, \$14.95, £11.95

"This revised and expanded edition of a clas-

sic [1962] textbook provides a concise introduction to basic concepts of acoustics and digital speech processing that are important to linguists, phoneticians, and speech scientists. The second edition includes four new chapters that cover new experimental techniques in acoustic phonetics made possible by the use of computers. Assuming no background in physics or mathematics, Ladefoged explains concepts that must be understood in using modern laboratory techniques for acoustic analysis, including resonances of the vocal tract and the relation of formants to different cavities; digital speech processing and computer storage of sound waves; and Fourier analysis and Linear Predictive Coding, the equations used most frequently in the analysis of speech sounds. Incorporating recent developments in our knowledge of the nature of speech, Ladefoged also updates the original edition's discussion of the basic properties of sound waves; variations in loudness, pitch, and quality of speech sounds; wave analysis; and the hearing and production of speech."—*From the publisher's announcement*