


The book was found

Speech And Language Processing, 2nd Edition



Synopsis

For undergraduate or advanced undergraduate courses in Classical Natural Language Processing, Statistical Natural Language Processing, Speech Recognition, Computational Linguistics, and Human Language Processing. An explosion of Web-based language techniques, merging of distinct fields, availability of phone-based dialogue systems, and much more make this an exciting time in speech and language processing. The first of its kind to thoroughly cover language technology at all levels and with all modern technologies this text takes an empirical approach to the subject, based on applying statistical and other machine-learning algorithms to large corporations. The authors cover areas that traditionally are taught in different courses, to describe a unified vision of speech and language processing. Emphasis is on practical applications and scientific evaluation. An accompanying Website contains teaching materials for instructors, with pointers to language processing resources on the Web. The Second Edition offers a significant amount of new and extended material. Supplements: Click on the "Resources" tab to View Downloadable Files: Solutions Power Point Lecture Slides - Chapters 1-5, 8-10, 12-13 and 24 Now Available! For additional resource visit the author website: <http://www.cs.colorado.edu/~martin/slp.html>

Book Information

Hardcover: 1024 pages

Publisher: Prentice Hall; 2nd edition (May 16, 2008)

Language: English

ISBN-10: 0131873210

ISBN-13: 978-0131873216

Product Dimensions: 7.4 x 1.7 x 9.4 inches

Shipping Weight: 3.4 pounds (View shipping rates and policies)

Average Customer Review: 4.1 out of 5 stars 24 customer reviews

Best Sellers Rank: #25,987 in Books (See Top 100 in Books) #1 in Books > Computers & Technology > Software > Voice Recognition #4 in Books > Computers & Technology > Computer Science > AI & Machine Learning > Natural Language Processing #6 in Books > Computers & Technology > Computer Science > AI & Machine Learning > Computer Vision & Pattern Recognition

Customer Reviews

Dan Jurafsky is an associate professor in the Department of Linguistics, and by courtesy in

Department of Computer Science, at Stanford University. Previously, he was on the faculty of the University of Colorado, Boulder, in the Linguistics and Computer Science departments and the Institute of Cognitive Science. He was born in Yonkers, New York, and received a B.A. in Linguistics in 1983 and a Ph.D. in Computer Science in 1992, both from the University of California at Berkeley. He received the National Science Foundation CAREER award in 1998 and the MacArthur Fellowship in 2002. He has published over 90 papers on a wide range of topics in speech and language processing.

James H. Martin is a professor in the Department of Computer Science and in the Department of Linguistics, and a fellow in the Institute of Cognitive Science at the University of Colorado at Boulder. He was born in New York City, received a B.S. in Computer Science from Columbia University in 1981 and a Ph.D. in Computer Science from the University of California at Berkeley in 1988. He has authored over 70 publications in computer science including the book *A Computational Model of Metaphor Interpretation*.

I give J&M five stars and they deserve it, and here's why. If you want learn to write natural language software, no other single book is as good as at least I've not found it. In fact, I bet they invented the genre. Pulling this together is not easy, and they do a creditable job. I know a lot more than I did before I read this book, and I've been writing linguistic software for over 30 years. As a linguist writing software (as opposed to the other way around), one can feel just a tad under siege these days. Google advertises that they don't have a single linguist on staff, and MS is ubiquitously quoted for saying that the quality of their software decreases for every linguist they hire. J&M, I'm happy to say, are above the fray. (What is supervised machine learning? Oh yeah, that's where your input was created by a linguist. Supervised or not, you're just playing number games on the foundation of a theoretical framework invented by linguists.) They provide a balanced account with historical perspective. I like them. They're cool. So on to picking nits... which is way more fun. What I really wanted is to read this book and then be able to sit down and write my own Python implementation of the forward/backward algorithm to train an HMM. I bobbed along through the book, perhaps experiencing a little bit of fuzziness around those probabilities, and came full stop at not quite right smack in the middle of my HMM forward/backward section. I'd done a practice run by training a neural net in Andrew Ng's machine learning course with Coursera. But I stared pretty hard for 3-4 hours at pages 189 and 190. And I mean I get it basically.

Alpha and beta represent the accumulated wisdom coming from the front and from the back. And then you take a kind of average to go from not quite ksi to ksi. But there are too many assumptions hidden in $P(X,Y|Z)/P(Y|Z)$. And this is an iterative algorithm, so how do you seed the counts? And I'm very annoyed by the phrase "denote the different conditioning of O". Okay, I can see the O is on the wrong side of the line. What does that mean? When I came to the next impasse, I didn't try as hard. It's already clear I'll have to go elsewhere for the silver bullet. (The next impasse, btw was the cepstrum "what do you mean you leave the graph the same and just replace the x-axis with something totally unrelated?" no Stanford professor, but what kind of math is that? I'm sure it means something to somebody, but not to me.) And drop the pseudo-code. If you're deadly serious about teaching me the HMM, then write out a working implementation in full in a real language like C or Python with the variables all initialized so I can copy and paste the code into my debugger and watch what happens to the numbers as I step through. I suspect J&M of compromising the pedagogical value of the book by deliberately withholding information from those brilliant Stanford students of theirs so they have something to quiz them on at the end of the chapter. But this is a mistake. Give us the answers. Give us all the answers. Give us the actual code for the HMM and then explain it. I will read the explanation, because my neck is on the line if my code blows up. There will still be plenty of questions left over for those students.

A wonderful book which is used in many Natural Language Processing courses. It covers a huge number of topics, and goes quite deeply into each of them. I didn't intend to purchase this book at first but when I realized how useful it would be to have a physical copy with me, I was not hesitate to get one.

While I'm only six or seven chapters into it, which are of an introductory nature, so far the book is excellent. For anyone interested in automated processing of natural speech, I think this would be a terrific addition. I'm learning a great deal from it, which is very much cementing my foundation in these concepts.

This is one of the books that I consider as a starting point / reference whenever I need to deal with a practical natural language processing (NLP) problem. I also have *Natural Language Processing*

with Python – on my shelf and it's wonderful in terms of providing a practical start for nearly any NLP problem but when the need arises to cover more ground both in terms of theory and practical pitfalls then Jurafsky & Martin is my guide. Natural language processing is a fast-moving target and it is impossible to know about the latest developments in the field without reading recent academic articles so nobody should expect to get the same information from this book, however mastering the concepts and algorithms in the book will provide the reader with the necessary background to understand state-of-the-art in NLP. Most of the exercises are very interesting but I wish they had some kind of difficulty level indicated next to them. Another criticism would be that more information on practical implementation details of the algorithms could have been given but I believe these minor criticisms does not lead to a four star rating. It is a very difficult project to give a comprehensive overview of the whole NLP field and Jurafsky & Martin achieved that.

I purchased this textbook initially for a class in natural language processing in the Biomedical Informatics domain. Throughout the semester, it provided itself as a excellent reference text and also an added bonus of providing problems that challenged me quite thoroughly. I would suggest this text as a must have if you are interested in the realm of natural language processing and speech processing.

This is a very helpful textbook which I used as part of an NLP class. It helped me understand the history, mathematics, and computer science of NLP work.

Needless to say, this is a classic in the NLP domain. It is different with most of other NLP book in that it focuses "real" computational linguistics but tons of other books focus on some toolkit or practical methodologies. The book is thorough and comprehensive and suitable for all levels of learners.

I love this book. It was easy to follow and a great read.

[Download to continue reading...](#)

Speech Therapy for Kids : Techniques and Parents Guide for Speech Therapy (speech therapy, speech therapy materials) Speech and Language Processing, 2nd Edition Discrete-Time Signal Processing (2nd Edition) (Prentice-Hall Signal Processing Series) The Complete Book of Speech Communication: A Workbook of Ideas and Activities for Students of Speech and Theatre The Passive Voice and Reported Speech: Your grammar torch to shed light on passive voice, reported

speech, complex subject, complex object and cleft (Brookgarbolt's treasure Book 2) Speech Science Primer: Physiology, Acoustics, and Perception of Speech Eulogy Writing: For Beginners! How To Write The Perfect Eulogy & Funeral Speech (Funeral, Obituary, Eulogy, Speech Writing, Public Speaking) Clinical Measurement of Speech & Voice (Speech Science) A Prosodic Model of Sign Language Phonology (Language, Speech, and Communication) Statistical Language Learning (Language, Speech, and Communication) Discrete-Time Signal Processing (3rd Edition) (Prentice-Hall Signal Processing Series) Natural Language Processing with Python: Analyzing Text with the Natural Language Toolkit Materials Processing: A Unified Approach to Processing of Metals, Ceramics and Polymers Multidimensional Digital Signal Processing (Prentice-Hall Signal Processing Series) Born to Talk: An Introduction to Speech and Language Development (6th Edition) A Guide to School Services in Speech-Language Pathology, Third Edition Here's How to Do Therapy: Hands on Core Skills in Speech-Language Pathology, Second Edition An Advanced Review of Speech-Language Pathology, 3rd Edition Speech to Print: Language Essentials for Teachers, Second Edition Sounds and Letters for Readers and Spellers: Phonemic Awareness Drills for Teachers and Speech-Language Pathologists

[Contact Us](#)

[DMCA](#)

[Privacy](#)

[FAQ & Help](#)