
Acces PDF Speech Recognition Algorithms Using Weighted Finite State Transducers Synthesis Lectures On Speech And Audio Processing

As recognized, adventure as well as experience approximately lesson, amusement, as competently as harmony can be gotten by just checking out a books **Speech Recognition Algorithms Using Weighted Finite State Transducers Synthesis Lectures On Speech And Audio Processing** furthermore it is not directly done, you could admit even more as regards this life, all but the world.

We provide you this proper as well as simple way to get those all. We manage to pay for Speech Recognition Algorithms Using Weighted Finite State Transducers Synthesis Lectures On Speech And Audio Processing and numerous books collections from fictions to scientific research in any way. among them is this Speech Recognition Algorithms Using Weighted Finite State Transducers Synthesis Lectures On Speech And Audio Processing that can be your partner.

KEY=AND - BRENDEN COHEN

SPEECH RECOGNITION ALGORITHMS BASED ON WEIGHTED FINITE-STATE TRANSDUCERS

Springer Nature This book introduces the theory, algorithms, and implementation techniques for efficient decoding in speech recognition mainly focusing on the Weighted Finite-State Transducer (WFST) approach. The decoding process for speech recognition is viewed as a search problem whose goal is to find a sequence of words that best matches an input speech signal. Since this process becomes computationally more expensive as the system vocabulary size increases, research has long been devoted to reducing the computational cost. Recently, the WFST approach has become an important state-of-the-art speech recognition technology, because it offers improved decoding speed with fewer recognition errors compared with conventional methods. However, it is not easy to understand all the algorithms used in this framework, and they are still in a black box for many people. In this book, we review the WFST approach and aim to provide comprehensive interpretations of WFST operations and decoding algorithms to help anyone who wants to understand, develop, and study WFST-based speech recognizers. We also mention recent advances in this framework and its applications to spoken language processing. Table of Contents: Introduction / Brief Overview of Speech Recognition / Introduction to Weighted Finite-State Transducers / Speech Recognition by Weighted Finite-State Transducers / Dynamic Decoders with On-the-fly WFST Operations / Summary and Perspective

FINITE-STATE TRANSDUCERS AND SPEECH RECOGNITION

"Finite-state automata and finite-state transducers have been extensively studied over the years. Recently, the theory of transducers has been generalized by Mohri for the weighted case. This generalization has allowed the use of finite-state transducers in a large variety of applications such as speech recognition. In this work, most of the algorithms for performing operations on weighted finite-state transducers are described in detail and analyzed. Then, an example of their use is given via a description of a speech recognition system based on them." --

WEIGHTED FINITE-STATE TRANSDUCERS IN SPEECH RECOGNITION [MICROFORM] : A COMPACTION ALGORITHM FOR NON-DETERMINIZABLE TRANSDUCERS

Montréal : Service des archives, Université de Montréal, Section Microfilm

BAYESIAN SPEECH AND LANGUAGE PROCESSING

Cambridge University Press A practical and comprehensive guide on how to apply Bayesian machine learning techniques to solve speech and language processing problems.

AUDIO SOURCE SEPARATION AND SPEECH ENHANCEMENT

John Wiley & Sons Learn the technology behind hearing aids, Siri, and Echo Audio source separation and speech enhancement aim to extract one or more source signals of interest from an audio recording involving several sound sources. These technologies are among the most studied in audio signal processing today and bear a critical role in the success of hearing aids, hands-free phones, voice command and other noise-robust audio analysis systems, and music post-production software. Research on this topic has followed three convergent paths, starting with sensor array processing, computational auditory scene analysis, and machine learning based approaches such as independent component analysis, respectively. This book is the first one to provide a comprehensive overview by presenting the common foundations and the differences between these techniques in a unified setting. Key features: Consolidated perspective on audio source separation and speech enhancement. Both historical perspective and latest advances in the field, e.g. deep neural networks. Diverse disciplines: array processing, machine learning, and statistical signal processing. Covers the most important techniques for both single-channel and multichannel processing. This book provides both introductory and advanced material suitable for people with basic knowledge of signal processing and machine learning. Thanks to its comprehensiveness, it will help students select a promising research track, researchers leverage the acquired cross-domain knowledge to design improved techniques, and engineers and

developers choose the right technology for their target application scenario. It will also be useful for practitioners from other fields (e.g., acoustics, multimedia, phonetics, and musicology) willing to exploit audio source separation or speech enhancement as pre-processing tools for their own needs.

WEIGHTED FINITE-STATE TRANSDUCERS IN SPEECH RECOGNITION

A COMPACTION ALGORITHM FOR NON-DETERMINIZABLE TRANSDUCERS

DFT-DOMAIN BASED SINGLE-MICROPHONE NOISE REDUCTION FOR SPEECH ENHANCEMENT

Springer Nature As speech processing devices like mobile phones, voice controlled devices, and hearing aids have increased in popularity, people expect them to work anywhere and at any time without user intervention. However, the presence of acoustical disturbances limits the use of these applications, degrades their performance, or causes the user difficulties in understanding the conversation or appreciating the device. A common way to reduce the effects of such disturbances is through the use of single-microphone noise reduction algorithms for speech enhancement. The field of single-microphone noise reduction for speech enhancement comprises a history of more than 30 years of research. In this survey, we wish to demonstrate the significant advances that have been made during the last decade in the field of discrete Fourier transform domain-based single-channel noise reduction for speech enhancement. Furthermore, our goal is to provide a concise description of a state-of-the-art speech enhancement system, and demonstrate the relative importance of the various building blocks of such a system. This allows the non-expert DSP practitioner to judge the relevance of each building block and to implement a close-to-optimal enhancement system for the particular application at hand. Table of Contents: Introduction / Single Channel Speech Enhancement: General Principles / DFT-Based Speech Enhancement Methods: Signal Model and Notation / Speech DFT Estimators / Speech Presence Probability Estimation / Noise PSD Estimation / Speech PSD Estimation / Performance Evaluation Methods / Simulation Experiments with Single-Channel Enhancement Systems / Future Directions

ACOUSTICAL IMPULSE RESPONSE FUNCTIONS OF MUSIC PERFORMANCE HALLS

Springer Nature Digital measurement of the analog acoustical parameters of a music performance hall is difficult. The aim of such work is to create a digital acoustical derivation that is an accurate numerical representation of the complex analog characteristics of the hall. The present study describes the exponential sine sweep (ESS) measurement process in the derivation of an acoustical impulse response function (AIRF) of three music performance halls in Canada. It examines specific difficulties of the process, such as preventing the external effects of the measurement transducers from corrupting the derivation, and provides solutions, such as the use of filtering techniques in order to remove such unwanted effects. In addition, the book presents a novel method of numerical verification through mean-squared error (MSE) analysis in order to determine how accurately the derived AIRF represents the acoustical behavior of the actual hall.

DEEP LEARNING FOR NLP AND SPEECH RECOGNITION

Springer This textbook explains Deep Learning Architecture, with applications to various NLP Tasks, including Document Classification, Machine Translation, Language Modeling, and Speech Recognition. With the widespread adoption of deep learning, natural language processing (NLP), and speech applications in many areas (including Finance, Healthcare, and Government) there is a growing need for one comprehensive resource that maps deep learning techniques to NLP and speech and provides insights into using the tools and libraries for real-world applications. Deep Learning for NLP and Speech Recognition explains recent deep learning methods applicable to NLP and speech, provides state-of-the-art approaches, and offers real-world case studies with code to provide hands-on experience. Many books focus on deep learning theory or deep learning for NLP-specific tasks while others are cookbooks for tools and libraries, but the constant flux of new algorithms, tools, frameworks, and libraries in a rapidly evolving landscape means that there are few available texts that offer the material in this book. The book is organized into three parts, aligning to different groups of readers and their expertise. The three parts are: Machine Learning, NLP, and Speech Introduction The first part has three chapters that introduce readers to the fields of NLP, speech recognition, deep learning and machine learning with basic theory and hands-on case studies using Python-based tools and libraries. Deep Learning Basics The five chapters in the second part introduce deep learning and various topics that are crucial for speech and text processing, including word embeddings, convolutional neural networks, recurrent neural networks and speech recognition basics. Theory, practical tips, state-of-the-art methods, experimentations and analysis in using the methods discussed in theory on real-world tasks. Advanced Deep Learning Techniques for Text and Speech The third part has five chapters that discuss the latest and cutting-edge research in the areas of deep learning that intersect with NLP and speech. Topics including attention mechanisms, memory augmented networks, transfer learning, multi-task learning, domain adaptation, reinforcement learning, and end-to-end deep learning for speech recognition are covered using case studies.

FINITE-STATE LANGUAGE PROCESSING

Language, Speech, and Communic Finite-state devices, which include finite-state automata, graphs, and finite-state transducers, are in wide use in many areas of computer science. Recently, there has been a resurgence of the use of finite-state devices in all aspects of computational linguistics, including dictionary encoding, text processing, and speech processing. This book describes the fundamental properties of finite-state devices and illustrates their uses.

Many of the contributors pioneered the use of finite-automata for different aspects of natural language processing. The topics, which range from the theoretical to the applied, include finite-state morphology, approximation of phrase-structure grammars, deterministic part-of-speech tagging, application of a finite-state intersection grammar, a finite-state transducer for extracting information from text, and speech recognition using weighted finite automata. The introduction presents the basic theoretical results in finite-state automata and transducers. These results and algorithms are described and illustrated with simple formal language examples as well as natural language examples. Contributors Douglas Appelt, John Bear, David Clemencau, Maurice Gross, Jerry R. Hobbs, David Israel, Megumi Kameyama, Lauri Karttunen, Kimmo Koskenniemi, Mehryar Mohri, Eric Laporte, Fernando C. N. Pereira, Michael D. Riley, Emmanuel Roche, Yves Schabes, Max D. Silberstein, Mark Stickel, Pasi Tapanainen, Mabry Tyson, Aro Voutilainen, Rebecca N. Wright

SPRINGER HANDBOOK OF SPEECH PROCESSING

Springer Science & Business Media This handbook plays a fundamental role in sustainable progress in speech research and development. With an accessible format and with accompanying DVD-Rom, it targets three categories of readers: graduate students, professors and active researchers in academia, and engineers in industry who need to understand or implement some specific algorithms for their speech-related products. It is a superb source of application-oriented, authoritative and comprehensive information about these technologies, this work combines the established knowledge derived from research in such fast evolving disciplines as Signal Processing and Communications, Acoustics, Computer Science and Linguistics.

COMPUTATIONAL PROCESSING OF THE PORTUGUESE LANGUAGE

7TH INTERNATIONAL WORKSHOP, PROPOR 2006, ITATIAIA, BRAZIL, MAY 13-17, 2006, PROCEEDINGS

Springer This book constitutes the thoroughly refereed proceedings of the 7th International Workshop on Computational Processing of the Portuguese Language, PROPOR 2006. The 20 revised full papers and 17 revised short papers presented here are organized in topical sections on automatic summarization, resources, translation, named entity recognition, tools and frameworks, systems and models, information extraction, speech processing, lexicon, morpho-syntactic studies, and Web, corpus and evaluation.

DEVELOPMENTS IN LANGUAGE THEORY

12TH INTERNATIONAL CONFERENCE, DLT 2008, KYOTO, JAPAN, SEPTEMBER 16-19, 2008, PROCEEDINGS

Springer Science & Business Media particular to Prof. Grzegorz Rozenberg for his valuable advice. This conference was dedicated to his 65th birthday. The conference was supported by Kyoto Sangyo University and the Japanese Society for the Promotion of Science.

ROBUSTNESS IN LANGUAGE AND SPEECH TECHNOLOGY

Springer Science & Business Media In this book we address robustness issues at the speech recognition and natural language parsing levels, with a focus on feature extraction and noise robust recognition, adaptive systems, language modeling, parsing, and natural language understanding. This book attempts to give a clear overview of the main technologies used in language and speech processing, along with an extensive bibliography to enable topics of interest to be pursued further. It also brings together speech and language technologies often considered separately. Robustness in Language and Speech Technology serves as a valuable reference and although not intended as a formal university textbook, contains some material that can be used for a course at the graduate or undergraduate level.

DETERMINIZATION OF STRING-TO-STRING/WEIGHT FINITE STATE TRANSUCERS IN SPEECH RECOGNITION

"The advantages and disadvantages of these algorithms have been discussed and analyzed. Special attention has been paid to the DSSW_determinization . The DSSW_determinization can be applied to the determinization of both determinizable and non-determinizable transducers, and it has a very low memory cost compared with that of the SSW_determinization and PSSW_determinization." --

AUTOMATIC SPEECH RECOGNITION ON MOBILE DEVICES AND OVER COMMUNICATION NETWORKS

Springer Science & Business Media The advances in computing and networking have sparked an enormous interest in deploying automatic speech recognition on mobile devices and over communication networks. This book brings together academic researchers and industrial practitioners to address the issues in this emerging realm and presents the reader with a comprehensive introduction to the subject of speech recognition in devices and networks. It covers network, distributed and embedded speech recognition systems.

COMPUTING AND COMBINATORICS

6TH ANNUAL INTERNATIONAL CONFERENCE, COCOON 2000, SYDNEY, AUSTRALIA, JULY 26-28, 2000 PROCEEDINGS

Springer This book constitutes the refereed proceedings of the 6th Annual International Conference on Computing and Combinatorics, COCOON 2000, held in Sydney, Australia in July 2000. The 44 revised full papers presented together

with two invited contributions were carefully reviewed and selected from a total of 81 submissions. The book offers topical sections on computational geometry; graph drawing; graph theory and algorithms; complexity, discrete mathematics, and number theory; online algorithms; parallel and distributed computing; combinatorial optimization; data structures and computational biology; learning and cryptography; and automata and quantum computing.

ADVANCES IN LOGIC BASED INTELLIGENT SYSTEMS

SELECTED PAPERS OF LAPTEC 2005

IOS Press LAPTEC2005 promoted the discussion and interaction between researchers and practitioners focused on both theoretical and practical disciplines concerning logics applied to technology, with diverse backgrounds including all kinds of intelligent systems having classical or non-classical logics as underlying common matters. It was the first time for LAPTEC to be held in a different country than Brazil since its birth in 2000, and this has made the congress more international. This book is dedicated to Emeritus Professor Atsuyuki Suzuki in commemoration of his honourable retirement from Shizuoka University, March 2005. Prof. Suzuki is learned in application of para consistent logic and has contributed many papers as a member of the program committee to LAPTEC since the beginning.

ADVANCES IN CHINESE SPOKEN LANGUAGE PROCESSING

HANDBOOK OF WEIGHTED AUTOMATA

Springer Science & Business Media The purpose of this Handbook is to highlight both theory and applications of weighted automata. Weighted finite automata are classical nondeterministic finite automata in which the transitions carry weights. These weights may model, e. g. , the cost involved when executing a transition, the amount of resources or time needed for this, or the probability or reliability of its successful execution. The behavior of weighted finite automata can then be considered as the function (suitably defined) associating with each word the weight of its execution. Clearly, weights can also be added to classical automata with infinite state sets like pushdown automata; this extension constitutes the general concept of weighted automata. To illustrate the diversity of weighted automata, let us consider the following scenarios. Assume that a quantitative system is modeled by a classical automaton in which the transitions carry as weights the amount of resources needed for their execution. Then the amount of resources needed for a path in this weighted automaton is obtained simply as the sum of the weights of its transitions. Given a word, we might be interested in the minimal amount of resources needed for its execution, i. e. , for the successful paths realizing the given word. In this example, we could also replace the “resources” by “profit” and then be interested in the maximal profit realized, correspondingly, by a given word.

FINITE-STATE TEXT PROCESSING

Springer Nature Weighted finite-state transducers (WFSTs) are commonly used by engineers and computational linguists for processing and generating speech and text. This book first provides a detailed introduction to this formalism. It then introduces Pynini, a Python library for compiling finite-state grammars and for combining, optimizing, applying, and searching finite-state transducers. This book illustrates this library's conventions and use with a series of case studies. These include the compilation and application of context-dependent rewrite rules, the construction of morphological analyzers and generators, and text generation and processing applications.

FORMAL LANGUAGES AND APPLICATIONS

Springer Formal Languages and Applications provides a comprehensive study-aid and self-tutorial for graduate students and researchers. The main results and techniques are presented in a readily accessible manner and accompanied by many references and directions for further research. This carefully edited monograph is intended to be the gateway to formal language theory and its applications, so it is very useful as a review and reference source of information in formal language theory.

ROBUSTNESS IN LANGUAGE AND SPEECH TECHNOLOGY

Springer Science & Business Media In this book we address robustness issues at the speech recognition and natural language parsing levels, with a focus on feature extraction and noise robust recognition, adaptive systems, language modeling, parsing, and natural language understanding. This book attempts to give a clear overview of the main technologies used in language and speech processing, along with an extensive bibliography to enable topics of interest to be pursued further. It also brings together speech and language technologies often considered separately. Robustness in Language and Speech Technology serves as a valuable reference and although not intended as a formal university textbook, contains some material that can be used for a course at the graduate or undergraduate level.

COMPUTATIONAL PROCESSING OF THE PORTUGUESE LANGUAGE

... INTERNATIONAL WORKSHOP, PROPOR ... : PROCEEDINGS

MACHINE TRANSLATION AND THE INFORMATION SOUP

THIRD CONFERENCE OF THE ASSOCIATION FOR MACHINE TRANSLATION IN THE AMERICAS, AMTA'98,

LANGHORNE, PA, USA, OCTOBER 28-31, 1998 PROCEEDINGS

Springer Machine Translation and the Information Soup! Over the past forty years, machine translation has grown from a tantalizing dream to a respectable and stable scientific-linguistic enterprise, with users, commercial systems, university research, and government participation. But until very recently, MT has been performed as a relatively distinct operation, somewhat isolated from other text processing. Today, this situation is changing rapidly. The explosive growth of the Web has brought multilingual text into the reach of nearly everyone with a computer. We live in a soup of information, an increasingly multilingual bouillabaisse. And to partake of this soup, we can use MT systems together with more and more tools and language processing technologies|information retrieval engines, -tomated text summarizers, and multimodal and multilingual displays. Though some of them may still be rather experimental, and though they may not quite fit together well yet, it is clear that the future will offer text manipulation systems that contain all these functions, seamlessly interconnected in various ways.

DISTANT SPEECH RECOGNITION

John Wiley & Sons A complete overview of distant automatic speech recognition The performance of conventional Automatic Speech Recognition (ASR) systems degrades dramatically as soon as the microphone is moved away from the mouth of the speaker. This is due to a broad variety of effects such as background noise, overlapping speech from other speakers, and reverberation. While traditional ASR systems underperform for speech captured with far-field sensors, there are a number of novel techniques within the recognition system as well as techniques developed in other areas of signal processing that can mitigate the deleterious effects of noise and reverberation, as well as separating speech from overlapping speakers. Distant Speech Recognition presents a contemporary and comprehensive description of both theoretic abstraction and practical issues inherent in the distant ASR problem. Key Features: Covers the entire topic of distant ASR and offers practical solutions to overcome the problems related to it Provides documentation and sample scripts to enable readers to construct state-of-the-art distant speech recognition systems Gives relevant background information in acoustics and filter techniques, Explains the extraction and enhancement of classification relevant speech features Describes maximum likelihood as well as discriminative parameter estimation, and maximum likelihood normalization techniques Discusses the use of multi-microphone configurations for speaker tracking and channel combination Presents several applications of the methods and technologies described in this book Accompanying website with open source software and tools to construct state-of-the-art distant speech recognition systems This reference will be an invaluable resource for researchers, developers, engineers and other professionals, as well as advanced students in speech technology, signal processing, acoustics, statistics and artificial intelligence fields.

LANGUAGE AND SPEECH PROCESSING

John Wiley & Sons Speech processing addresses various scientific and technological areas. It includes speech analysis and variable rate coding, in order to store or transmit speech. It also covers speech synthesis, especially from text, speech recognition, including speaker and language identification, and spoken language understanding. This book covers the following topics: how to realize speech production and perception systems, how to synthesize and understand speech using state-of-the-art methods in signal processing, pattern recognition, stochastic modelling computational linguistics and human factor studies.

SIAM JOURNAL ON COMPUTING

AUTOMATIC RECOGNITION OF DYSARTHIC SPEECH

LAP Lambert Academic Publishing Dysarthria is a motor speech disorder characterized by weakness, paralysis, or poor coordination of the muscles responsible for speech. Although automatic speech recognition (ASR) systems have been developed for disordered speech, factors such as low intelligibility and limited phonemic repertoire decrease speech recognition accuracy. Furthermore, conventional speaker adaptation algorithms that improve normal speech recognition may not perform as well on dysarthric speakers. Instead of adapting the system, two main techniques are proposed to model the pronunciation errors made by the speaker: (1) a set of discrete hidden Markov models (termed as "metamodels") that incorporate a model of the speaker's phonetic confusion-matrix into the ASR process; and (2) a network of Weighted Finite-State Transducers (WFSTs) at the confusion-matrix, word and language levels. These error modelling techniques attempt to correct the errors made at the phonetic level and make use of a language model to find the best estimate of the correct word sequence. Hence, these techniques when integrated into the speech recognition process performed significant error correction and improved speech recognition

CONNECTIONIST SPEECH RECOGNITION

A HYBRID APPROACH

Springer Science & Business Media Connectionist Speech Recognition: A Hybrid Approach describes the theory and implementation of a method to incorporate neural network approaches into state of the art continuous speech recognition systems based on hidden Markov models (HMMs) to improve their performance. In this framework, neural networks (and in particular, multilayer perceptrons or MLPs) have been restricted to well-defined subtasks of the whole system, i.e. HMM emission probability estimation and feature extraction. The book describes a successful five-year international collaboration between the authors. The lessons learned form a case study that demonstrates how

hybrid systems can be developed to combine neural networks with more traditional statistical approaches. The book illustrates both the advantages and limitations of neural networks in the framework of a statistical systems. Using standard databases and comparison with some conventional approaches, it is shown that MLP probability estimation can improve recognition performance. Other approaches are discussed, though there is no such unequivocal experimental result for these methods. Connectionist Speech Recognition is of use to anyone intending to use neural networks for speech recognition or within the framework provided by an existing successful statistical approach. This includes research and development groups working in the field of speech recognition, both with standard and neural network approaches, as well as other pattern recognition and/or neural network researchers. The book is also suitable as a text for advanced courses on neural networks or speech processing.

SPEECH & LANGUAGE PROCESSING

Pearson Education India

AUTOMATIC SPEECH AND SPEAKER RECOGNITION

LARGE MARGIN AND KERNEL METHODS

John Wiley & Sons This book discusses large margin and kernel methods for speech and speaker recognition Speech and Speaker Recognition: Large Margin and Kernel Methods is a collation of research in the recent advances in large margin and kernel methods, as applied to the field of speech and speaker recognition. It presents theoretical and practical foundations of these methods, from support vector machines to large margin methods for structured learning. It also provides examples of large margin based acoustic modelling for continuous speech recognizers, where the grounds for practical large margin sequence learning are set. Large margin methods for discriminative language modelling and text independent speaker verification are also addressed in this book. Key Features: Provides an up-to-date snapshot of the current state of research in this field Covers important aspects of extending the binary support vector machine to speech and speaker recognition applications Discusses large margin and kernel method algorithms for sequence prediction required for acoustic modeling Reviews past and present work on discriminative training of language models, and describes different large margin algorithms for the application of part-of-speech tagging Surveys recent work on the use of kernel approaches to text-independent speaker verification, and introduces the main concepts and algorithms Surveys recent work on kernel approaches to learning a similarity matrix from data This book will be of interest to researchers, practitioners, engineers, and scientists in speech processing and machine learning fields.

LANGUAGE ENGINEERING FOR LESSER-STUDIED LANGUAGES

INTELLIGENT SIGNAL PROCESSING

Wiley-IEEE Press Intelligent signal processing (ISP) differs fundamentally from the classical approach to statistical signal processing in that the input-output behavior of a complex system is modeled by using an artificial intelligence capable of optimizing results.

2020 54TH ASILOMAR CONFERENCE ON SIGNALS, SYSTEMS, AND COMPUTERS

The Asilomar Conference provides an informal venue for technical exchange in the areas of signal processing, communication, system theory, biomedical signal processing, array processing, and computer architecture and arithmetic

PROCEEDINGS OF THE 1998 IEEE INTERNATIONAL CONFERENCE ON ACOUSTICS, SPEECH, AND SIGNAL PROCESSING

MAY 12-15, 1998, WASHINGTON STATE CONVENTION AND TRADE CENTER, SEATTLE, WASHINGTON (USA)

MATHEMATICAL REVIEWS

PROCEEDINGS

MAY 12-15, 1998, WASHINGTON STATE CONVENTION AND TRADE CENTER, SEATTLE, WA (USA).. COMMUNICATIONS SIGNAL PROCESSING. 1998,6

LINGUISTICS AND LANGUAGE BEHAVIOR ABSTRACTS

LLBA.

40TH ANNUAL MEETING OF THE ASSOCIATION FOR COMPUTATIONAL LINGUISTICS

7-12 JULY 2002, UNIVERSITY OF PENNSYLVANIA, PHILADELPHIA, PENNSYLVANIA, USA.

Morgan Kaufmann