
Download Ebook Automatic Speech Recognition On Mobile Devices And Over Communication Networks Advances In Computer Vision And Pattern Recognition

Recognizing the pretentiousness ways to get this books **Automatic Speech Recognition On Mobile Devices And Over Communication Networks Advances In Computer Vision And Pattern Recognition** is additionally useful. You have remained in right site to start getting this info. get the Automatic Speech Recognition On Mobile Devices And Over Communication Networks Advances In Computer Vision And Pattern Recognition connect that we present here and check out the link.

You could purchase lead Automatic Speech Recognition On Mobile Devices And Over Communication Networks Advances In Computer Vision And Pattern Recognition or get it as soon as feasible. You could speedily download this Automatic Speech Recognition On Mobile Devices And Over Communication Networks Advances In Computer Vision And Pattern Recognition after getting deal. So, once you require the books swiftly, you can straight acquire it. Its fittingly categorically easy and consequently fats, isnt it? You have to favor to in this ventilate

703RJV - ANGEL BRIANA

"This book introduces the readers to the various aspects of visual speech recognitions, including lip segmentation from video sequence, lip feature extraction and modeling, feature fusion and classifier design for visual speech recognition and speaker verification" résumé de l'éditeur.

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. 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.

Robust Automatic Speech Recognition: A Bridge to Practical Applications establishes a solid foundation for automatic speech recognition that is robust against acoustic environmental distortion. It provides a thorough overview of classical and modern noise-and reverberation robust techniques that have been developed over the past thirty years, with an emphasis on practical methods that have been proven to be successful and which are likely to be further developed for future applications. The strengths and weaknesses of robustness-enhancing speech recognition techniques are carefully analyzed. The book covers noise-robust techniques designed for acous-

tic models which are based on both Gaussian mixture models and deep neural networks. In addition, a guide to selecting the best methods for practical applications is provided. The reader will: Gain a unified, deep and systematic understanding of the state-of-the-art technologies for robust speech recognition Learn the links and relationship between alternative technologies for robust speech recognition Be able to use the technology analysis and categorization detailed in the book to guide future technology development Be able to develop new noise-robust methods in the current era of deep learning for acoustic modeling in speech recognition The first book that provides a comprehensive review on noise and reverberation robust speech recognition methods in the era of deep neural networks Connects robust speech recognition techniques to machine learning paradigms with rigorous mathematical treatment Provides elegant and structural ways to categorize and analyze noise-robust speech recognition techniques Written by leading researchers who have been actively working on the subject matter in both industrial

and academic organizations for many years

Abstract: "Automatic speech recognition (ASR) is a computerized speech-to-text process, in which speech is usually recorded with acoustical microphones by capturing air pressure changes. This kind of air-transmitted speech signal is prone to two kinds of problems related to noise robustness and applicability. The former means the mixing of speech signal and ambient noise usually deteriorates ASR performance. The latter means speech could be overheard easily on the air-transmission channel, and this often results in privacy loss or annoyance to other people. This thesis research solves these two problems by using channels that contact the human body without air transmission, i.e., by vibrocervigraphic and electromyographic methods. The vibrocervigraphic (VCG) method measures the throat vibration with a ceramic piezoelectric transducer contact to the skin on the neck, and the electromyographic (EMG) method measures the muscular electric potential with a set of electrodes attached to the skin where the articulatory

muscles underlie. The VCG and EMG methods are inherently more robust to ambient noise, and they make it possible to recognize whispered and silent speech to improve applicability. The major contribution of this dissertation includes feature design and adaptation for optimizing features, acoustic model adaptation for adapting traditional acoustic models onto different feature spaces, and articulatory feature classification for incorporating articulatory information to improve recognition. For VCG ASR, the combination of feature transformation methods and maximum a posteriori adaptation improves the recognition accuracy even with a very small data set. On top of that, additive performance gain is achieved by applying maximum likelihood linear regression and feature space adaptation with different data granularities in order to adapt to channel variations as well as to speaker variations. For EMG ASR, we propose the Concise EMG feature that extracts representative EMG characteristics. It improves the recognition accuracy and advances the EMG ASR research from isolated word recognition to phone-

based continuous speech recognition. Articulatory features are studied in both VCG and EMG ASR to analyze the systems and improve recognition accuracy. These techniques are demonstrated to be effective on both experimental evaluations and prototype applications." Automatic Speech Recognition (ASR) is one of the most important applications in the area of cognitive computing. Fast and accurate ASR is emerging as a key application for mobile and wearable devices. These devices, such as smartphones, have incorporated speech recognition as one of the main interfaces for user interaction. This trend towards voice-based user interfaces is likely to continue in the next years which is changing the way of human-machine interaction. Effective speech recognition systems require real-time recognition, which is challenging for mobile devices due to the compute-intensive nature of the problem and the power constraints of such systems and involves a huge effort for CPU architectures to reach it. GPU architectures offer parallelization capabilities which can be exploited to increase the performance of speech recognition sys-

tems. However, efficiently utilizing the GPU resources for speech recognition is also challenging, as the software implementations exhibit irregular and unpredictable memory accesses and poor temporal locality. The purpose of this thesis is to study the characteristics of ASR systems running on low-power mobile devices in order to propose different techniques to improve performance and energy consumption. We propose several software-level optimizations driven by the power/performance analysis. Unlike previous proposals that trade accuracy for performance by reducing the number of Gaussians evaluated, we maintain accuracy and improve performance by effectively using the underlying CPU microarchitecture. We use a refactored implementation of the G03 evaluation code to ameliorate the impact of branches. Then, we exploit the vector unit available on most modern CPUs to boost G03 computation, introducing a novel memory layout for storing the means and variances of the Gaussians in order to maximize the effectiveness of vectorization. In addition, we compute the Gaussians for multiple frames in parallel, signifi-

cantly reducing memory bandwidth usage. Our experimental results show that the proposed optimizations provide 2.68x speedup over the baseline Pocketsphinx decoder on a high-end Intel Skylake CPU, while achieving 61% energy savings. On a modern ARM Cortex-A57 mobile processor our techniques improve performance by 1.85x, while providing 59% energy savings without any loss in the accuracy of the ASR system. Secondly, we propose a register renaming technique that exploits register reuse to reduce the pressure on the register file. Our technique leverages physical register sharing by introducing minor changes in the register map table and the issue queue. We evaluated our renaming technique on top of a modern out-of-order processor. The proposed scheme supports precise exceptions and we show that it results in 9.5% performance improvements for G03 evaluation. Our experimental results show that the proposed register renaming scheme provides 6% speedup on average for the SPEC2006 benchmarks. Alternatively, our renaming scheme achieves the same performance while reducing the

number of physical registers by 10.5%. Finally, we propose a hardware accelerator for G03 evaluation that reduces the energy consumption by three orders of magnitude compared to solutions based on CPUs and GPUs. The proposed accelerator implements a lazy evaluation scheme where Gaussians are computed on demand, avoiding 50% of the computations. Furthermore, it employs a novel clustering scheme to reduce the size of the G03 parameters, which results in 8x memory bandwidth savings with a negligible impact on accuracy. Finally, it includes a novel memoization scheme that avoids 74.88% of floating-point operations. The end design provides a 164x speedup and 3532x energy reduction when compared with a highly-tuned implementation running on a modern mobile CPU. Compared to a state-of-the-art mobile GPU, the G03 accelerator achieves 5.89x speedup over a highly optimized CUDA implementation, while reducing energy by 241x.

Foreword Looking back the past 30 years. we have seen steady progress made in the area of speech science and technology. I still remember the excitement in the

late seventies when Texas Instruments came up with a toy named "Speak-and-Spell" which was based on a VLSI chip containing the state-of-the-art linear prediction synthesizer. This caused a speech technology fever among the electronics industry. Particularly, applications of automatic speech recognition were rigorously attempted by many companies, some of which were start-ups founded just for this purpose. Unfortunately, it did not take long before they realized that automatic speech recognition technology was not mature enough to satisfy the need of customers. The fever gradually faded away. In the meantime, constant efforts have been made by many researchers and engineers to improve the automatic speech recognition technology. Hardware capabilities have advanced impressively since that time. In the past few years, we have been witnessing and experiencing the advent of the "Information Revolution." What might be called the second surge of interest to commercialize speech technology as a natural interface for man-machine communication began in much better shape than the first one. With computers much

more powerful and faster, many applications look realistic this time. However, there are still tremendous practical issues to be overcome in order for speech to be truly the most natural interface between humans and machines.

The term Intelligent Environments (IEs) refers to physical spaces in which IT and other pervasive computing technologies are combined and used to achieve specific goals for the user, the environment, or both. The ultimate objective of IEs is to enrich user experience, improve management of the environment in question and increase user awareness. This book presents the proceedings of the following workshops, which formed part of the 12th International Conference on Intelligent Environments (IE16), held in London, UK, in September 2016: the 5th International Workshop on Smart Offices and Other Workplaces (SOOW'16); the 5th International Workshop on the Reliability of Intelligent Environments (WoRIE'16); the 1st International Workshop on Legal Issues in Intelligent Environments (LIIE'2016); the 2nd International Symposium on Future Intelligent Educational Environments and Learning (SO-

FIEE'16); the 2nd International Workshop on Future Internet and Smart Networks (FI&SN'2016); the International Workshop on Intelligent Environments Supporting Healthcare and Well-being (WISHWell'2016); the International Workshop on Computation Sustainability, Technologies and Applications (CoS-TA'2016); the Creative Science 2016 (CS'16) and Cloud-of-Things 2016 (CoT'16); the Workshop on Wireless Body Area Networks for Personal Monitoring in Intelligent Environments (WBAN-PMIE); and the Physical Computing Workshop. The workshops focused on the development of advanced intelligent environments, as well as newly emerging and rapidly evolving topics, emphasizing the multi-disciplinary and transversal aspects of IEs, as well as cutting-edge topics. The book will be of interest to all those whose work involves them in the use of intelligent environments.

The book presents current research and developments in multilingual speech recognition. The author presents a Multilingual Phone Recognition System (Multi-PRS), developed using a common multilingual phone-set de-

rived from the International Phonetic Alphabets (IPA) based transcription of six Indian languages - Kannada, Telugu, Bengali, Odia, Urdu, and Assamese. The author shows how the performance of Multi-PRS can be improved using tandem features. The book compares Monolingual Phone Recognition Systems (Mono-PRS) versus Multi-PRS and baseline versus tandem system. Methods are proposed to predict Articulatory Features (AFs) from spectral features using Deep Neural Networks (DNN). Multi-task learning is explored to improve the prediction accuracy of AFs. Then, the AFs are explored to improve the performance of Multi-PRS using lattice rescoring method of combination and tandem method of combination. The author goes on to develop and evaluate the Language Identification followed by Monolingual phone recognition (LID--Mono) and common multilingual phone-set based multilingual phone recognition systems.

Automatic speech recognition and speaker recognition have a lot of applications in personal identification, access control and in the new man-machine-interface paradigm. The ex-

isting applications in voice-activated embedded systems solve the problem of recognition of the spoken words only or the problem of recognition of a speaker through the words uttered only. The goal of this project, therefore, is the development of a robust algorithm for both speech recognition and speaker verification. An example of a target application of this work is speech dialing of mobile phones with a speaker verification front-end in order to effect access control. In view of the memory and computational constraints of embedded systems, the dynamic time warping algorithm is used. This project only considers isolated spoken digits. The developed algorithm is coded in C language and can be ported to firmware for Arabic numeral digit recognition with a speaker verification front end for an embedded system like mobile phones. The system produced a FAR of 13.33% and a FRR of 24.3% for a total of 70 true claims and 30 false claims. It also had a word accuracy of 96.7%. This book covers language modeling and automatic speech recognition for inflective languages (e.g. Slavic languages), which represent roughly

half of the languages spoken in Europe. These languages do not perform as well as English in speech recognition systems and it is therefore harder to develop an application with sufficient quality for the end user. The authors describe the most important language features for the development of a speech recognition system. This is then presented through the analysis of errors in the system and the development of language models and their inclusion in speech recognition systems, which specifically address the errors that are relevant for targeted applications. The error analysis is done with regard to morphological characteristics of the word in the recognized sentences. The book is oriented towards speech recognition with large vocabularies and continuous and even spontaneous speech. Today such applications work with a rather small number of languages compared to the number of spoken languages.

Two Top Industry Leaders Speak Out Judith Markowitz When Amy asked me to co-author the foreword to her new book on advances in speech recognition, I was honored. Amy's work has al-

ways been infused with creative intensity, so I knew the book would be as interesting for established speech professionals as for readers new to the speech-processing industry. The fact that I would be writing the foreword with Bill Scholz made the job even more enjoyable. Bill and I have known each other since he was at UNISYS directing projects that had a profound impact on speech-recognition tools and applications. Bill Scholz The opportunity to prepare this foreword with Judith provides me with a rare opportunity to collaborate with a seasoned speech professional to identify numerous significant contributions to the field offered by the contributors whom Amy has recruited. Judith and I have had our eyes opened by the ideas and analyses offered by this collection of authors. Speech recognition no longer needs be relegated to the category of an experimental future technology; it is here today with sufficient capability to address the most challenging of tasks. And the point-click-type approach to GUI control is no longer sufficient, especially in the context of limitations of modern day hand held devices. Instead, VUI and

GUI are being integrated into unified multimodal solutions that are maturing into the fundamental paradigm for computer-human interaction in the future.

This book provides a comprehensive overview of the recent advancement in the field of automatic speech recognition with a focus on deep learning models including deep neural networks and many of their variants. This is the first automatic speech recognition book dedicated to the deep learning approach. In addition to the rigorous mathematical treatment of the subject, the book also presents insights and theoretical foundation of a series of highly successful deep learning models.

This book focuses primarily on speech recognition and the related tasks such as speech enhancement and modeling. This book comprises 3 sections and thirteen chapters written by eminent researchers from USA, Brazil, Australia, Saudi Arabia, Japan, Ireland, Taiwan, Mexico, Slovakia and India. Section 1 on speech recognition consists of seven chapters. Sections 2 and 3 on speech enhancement and speech modeling have three chapters each

respectively to supplement section 1. We sincerely believe that thorough reading of these thirteen chapters will provide comprehensive knowledge on modern speech recognition approaches to the readers.

Voice user interfaces (VUIs) are becoming all the rage today. But how do you build one that people can actually converse with? Whether you're designing a mobile app, a toy, or a device such as a home assistant, this practical book guides you through basic VUI design principles, helps you choose the right speech recognition engine, and shows you how to measure your VUI's performance and improve upon it. Author Cathy Pearl also takes product managers, UX designers, and VUI designers into advanced design topics that will help make your VUI not just functional, but great. Understand key VUI design concepts, including command-and-control and conversational systems. Decide if you should use an avatar or other visual representation with your VUI. Explore speech recognition technology and its impact on your design. Take your VUI above and beyond the basic exchange of information. Learn practical

cal ways to test your VUI application with users. Monitor your app and learn how to quickly improve performance. Get real-world examples of VUIs for home assistants, smartwatches, and car systems.

Speech recognition technology is being increasingly employed in human-machine interfaces. A remaining problem however is the robustness of this technology to non-native accents, which still cause considerable difficulties for current systems. In this book, methods to overcome this problem are described. A speaker adaptation algorithm that is capable of adapting to the current speaker with just a few words of speaker-specific data based on the MLLR principle is developed and combined with confidence measures that focus on phone durations as well as on acoustic features. Furthermore, a specific pronunciation modelling technique that allows the automatic derivation of non-native pronunciations without using non-native data is described and combined with the previous techniques to produce a robust adaptation to non-native accents in an automatic speech recognition system.

Automatic speech recognition (ASR) is a very attractive means for human-machine interaction. The degree of maturity reached by speech recognition technologies during recent years allows the development of applications that use them. In particular, ASR shows an enormous potential in mobile environments, where devices such as mobile phones or PDAs are used, and for Internet Protocol (IP) applications. *Speech Recognition Over Digital Channels* is the first book of its kind to offer a complete system comprehension, addressing the topics of distributed and network-based speech recognition issues and standards, the concepts of speech processing and transmission, and system architectures and robustness. Describes the different client/server architectures for remote speech recognition systems, by means of which the client transmits speech parameters through a digital channel to a remote recognition server. Focuses on robustness against both adverse acoustic environments (in the front-end) and bit errors/packet loss. Discusses four ETSI standards for distributed speech recognition; the understanding of the stan-

dards and the technologies behind them. Provides the necessary background for the comprehension of remote speech recognition technologies. This book will appeal to a wide-ranging audience: engineers using speech recognition systems, researchers involved in ASR systems and those interested in processing and transmitting speech such as signal processing and communications communities. It will also be of interest to technical experts requiring an understanding of recognition over mobile and IP networks, and post-graduate students working on robust speech processing.

"*Mobile Speech and Advanced Natural Language Solutions*" presents the discussion of the most recent advances in intelligent human-computer interaction, including fascinating new study findings on talk-in-interaction, which is the province of conversation analysis, a subfield in sociology/sociolinguistics, a new and emerging area in natural language understanding. Editors Amy Neustein and Judith A. Markowitz have recruited a talented group of contributors to introduce the next generation natural language technologies for practical speech

processing applications that serve the consumer's need for well-functioning natural language-driven personal assistants and other mobile devices, while also addressing business' need for better functioning IVR-driven call centers that yield a more satisfying experience for the caller. This anthology is aimed at two distinct audiences: one consisting of speech engineers and system developers; the other comprised of linguists and cognitive scientists. The text builds on the experience and knowledge of each of these audiences by exposing them to the work of the other.

Chapters in the first part of the book cover all the essential speech processing techniques for building robust, automatic speech recognition systems: the representation for speech signals and the methods for speech-features extraction, acoustic and language modeling, efficient algorithms for searching the hypothesis space, and multimodal approaches to speech recognition. The last part of the book is devoted to other speech processing applications that can use the information from automatic speech recognition for speaker identification and tracking, for prosody mod-

eling in emotion-detection systems and in other speech processing applications that are able to operate in real-world environments, like mobile communication services and smart homes.

Speech Recognition has a long history of being one of the difficult problems in Artificial Intelligence and Computer Science. As one goes from problem solving tasks such as puzzles and chess to perceptual tasks such as speech and vision, the problem characteristics change dramatically: knowledge poor to knowledge rich; low data rates to high data rates; slow response time (minutes to hours) to instantaneous response time. These characteristics taken together increase the computational complexity of the problem by several orders of magnitude. Further, speech provides a challenging task domain which embodies many of the requirements of intelligent behavior: operate in real time; exploit vast amounts of knowledge, tolerate errorful, unexpected unknown input; use symbols and abstractions; communicate in natural language and learn from the environment. Voice input to computers offers a number of advantages. It provides a natural, fast,

hands free, eyes free, location free input medium. However, there are many as yet unsolved problems that prevent routine use of speech as an input device by non-experts. These include cost, real time response, speaker independence, robustness to variations such as noise, microphone, speech rate and loudness, and the ability to handle non-grammatical speech. Satisfactory solutions to each of these problems can be expected within the next decade. Recognition of unrestricted spontaneous continuous speech appears unsolvable at present. However, by the addition of simple constraints, such as clarification dialog to resolve ambiguity, we believe it will be possible to develop systems capable of accepting very large vocabulary continuous speechdictation.

This Edited Volume gathers a selection of refereed and revised papers originally presented at the Third International Symposium on Signal Processing and Intelligent Recognition Systems (SIRS'17), held on September 13-16, 2017 in Manipal, India. The papers offer stimulating insights into biometrics, digital watermarking, recognition systems, im-

age and video processing, signal and speech processing, pattern recognition, machine learning and knowledge-based systems. Taken together, they offer a valuable resource for all researchers and scientists engaged in the various fields of signal processing and related areas.

Over the past few years, speech recognition technology performance on tasks ranging from isolated digit recognition to conversational speech has dramatically improved. Performance on limited recognition tasks in noise-free environments is comparable to that achieved by human transcribers. This advancement in automatic speech recognition technology along with an increase in the compute power of mobile devices, standardization of communication protocols, and the explosion in the popularity of the mobile devices, has created an interest in flexible voice interfaces for mobile devices. However, speech recognition performance degrades dramatically in mobile environments which are inherently noisy. In the recent past, a great amount of effort has been spent on the development of front ends based on advanced noise robust ap-

proaches. The primary objective of this thesis was to analyze the performance of two advanced front ends, referred to as the QIO and MFA front ends, on a speech recognition task based on the Wall Street Journal database. Though the advanced front ends are shown to achieve a significant improvement over an industry-standard baseline front end, this improvement is not operationally significant. Further, we show that the results of this evaluation were not significantly impacted by suboptimal recognition system parameter settings. Without any front end-specific tuning, the MFA front end outperforms the QIO front end by 9.6% relative. With tuning, the relative performance gap increases to 15.8%. Finally, we also show that mismatched microphone and additive noise evaluation conditions resulted in a significant degradation in performance for both front ends.

Automatic speech recognition suffers from a lack of robustness with respect to noise, reverberation and interfering speech. The growing field of speech recognition in the presence of missing or uncertain input data seeks to

ameliorate those problems by using not only a preprocessed speech signal but also an estimate of its reliability to selectively focus on those segments and features that are most reliable for recognition. This book presents the state of the art in recognition in the presence of uncertainty, offering examples that utilize uncertainty information for noise robustness, reverberation robustness, simultaneous recognition of multiple speech signals, and audiovisual speech recognition. The book is appropriate for scientists and researchers in the field of speech recognition who will find an overview of the state of the art in robust speech recognition, professionals working in speech recognition who will find strategies for improving recognition results in various conditions of mismatch, and lecturers of advanced courses on speech processing or speech recognition who will find a reference and a comprehensive introduction to the field. The book assumes an understanding of the fundamentals of speech recognition using Hidden Markov Models.

The portable device and mobile phone market has witnessed rapid growth in the last few years with the

emergence of several revolutionary products such as mobile TV, converging iPhone and digital cameras that combine music, phone and video functionalities into one device. The proliferation of this market has further benefited from the competition in software and applications for smart phones such as Google's Android operating system and Apple's iPhone App-Store, stimulating tens of thousands of mobile applications that are made available by individual and enterprise developers. Whereas the mobile device has become ubiquitous in people's daily life not only as a cellular phone but also as a media player, a mobile computing device, and a personal assistant, it is particularly important to address challenges timely in applying advanced pattern recognition, signal, information and multimedia processing techniques, and new emerging networking technologies to such mobile systems. The primary objective of this book is to foster interdisciplinary discussions and research in mobile multimedia processing techniques, applications and systems, as well as to provide stimulus to researchers on pushing the frontier of emerg-

ing new technologies and applications. One attempt on such discussions was the organization of the First International Workshop of Mobile Multimedia Processing (WMMP 2008), held in Tampa, Florida, USA, on December 7, 2008. About 30 papers were submitted from 10 countries across the USA, Asia and Europe.

Research is summarized on problems concerned with the interpretation of the acoustic parameters of speech and with the use of linguistic information in processing the discrete code derived from these parameters. The summary discusses work on: the theory of phonology, including development of a phonetic theory based on physiological parameters, and a phonemic theory in natural language and mathematical forms; procedures for converting acoustic parameters to phone types or sets of phone types that can be specified on a phonetic basis; a structural description of the phonology of midwest American English dialect; and lexical procedures for automatic speech recognition. (Author).

This book provides a cross-disciplinary reference to speech in mobile

and pervasive environments. *Speech in Mobile and Pervasive Environments* addresses the issues related to speech processing on resource-constrained mobile devices. These include speech recognition in noisy environments, specialised hardware for speech recognition and synthesis, the use of context to enhance recognition and user experience, and the emerging software standards required for interoperability. This book takes a multidisciplinary look at these matters, while offering an insight into the opportunities and challenges of speech processing in mobile environs. In developing regions, speech-on-mobile is set to play a momentous role, socially and economically; the authors discuss how voice-based solutions and applications offer a compelling and natural solution in this setting. *Key Features* Provides a holistic overview of all speech technology related topics in the context of mobility. *Brings together* the latest research in a logically connected way in a single volume. *Covers* hardware, embedded recognition and synthesis, distributed speech recognition, software technologies, contextual interfaces. *Discusses* multimodal dia-

logue systems and their evaluation. Introduces speech in mobile and pervasive environments for developing regions. This book provides a comprehensive overview for beginners and experts alike. It can be used as a textbook for advanced undergraduate and postgraduate students in electrical engineering and computer science. Students, practitioners or researchers in the areas of mobile computing, speech processing, voice applications, human-computer interfaces, and information and communication technologies will also find this reference insightful. For experts in the above domains, this book complements their strengths. In addition, the book will serve as a guide to practitioners working in telecom-related industries. Automatic speech recognition (ASR) systems are finding increasing use in everyday life. Many of the commonplace environments where the systems are used are noisy, for example users calling up a voice search system from a busy cafeteria or a street. This can result in degraded speech recordings and adversely affect the performance of speech recognition systems. As the use of ASR systems in-

creases, knowledge of the state-of-the-art techniques to deal with such problems becomes critical to system and application engineers and researchers who work with or on ASR technologies. This book presents a comprehensive survey of the state-of-the-art in techniques used to improve the robustness of speech recognition systems to these degrading external influences. Key features: Reviews all the main noise robust ASR approaches, including signal separation, voice activity detection, robust feature extraction, model compensation and adaptation, missing data techniques and recognition of reverberant speech. Acts as a timely exposition of the topic in light of more widespread use in the future of ASR technology in challenging environments. Addresses robustness issues and signal degradation which are both key requirements for practitioners of ASR. Includes contributions from top ASR researchers from leading research units in the field. Learning vocabulary using ASR is a unique concept since it demands language learners to speak up in order to find word definitions. With the advancement of technology,

it has become possible to incorporate ASR to computers, mobile phones, and other smart digital devices to learn L2 vocabulary. Two groups of adult, ESOL learners were given a vocabulary activity which had to be completed in class by using Google and Apple based ASR assistants in their mobile phones. Then, they were given a brief questionnaire inquiring about their perceptions on using ASR in the classroom. After the study, the teachers of the two classes were also asked about their perceptions about using ASR in the classroom. The data gathered from students' questionnaire was analyzed under three qualitative evaluation criteria. Finally, the students' perceptions about ASR were compared with teachers' perceptions. While students cited many advantages of using ASR in the classroom to learn vocabulary, overall they claimed that mispronunciation of words and verbal commands became a hindrance in learning L2 vocabulary via ASR. On the contrary, the teachers thought that ASR helped student to improve their pronunciation. Thus, there was also a mismatch between students' perceptions and teachers' per-

ceptions.--

Spoken language understanding (SLU) is an emerging field in between speech and language processing, investigating human/ machine and human/ human communication by leveraging technologies from signal processing, pattern recognition, machine learning and artificial intelligence. SLU systems are designed to extract the meaning from speech utterances and its applications are vast, from voice search in mobile devices to meeting summarization, attracting interest from both commercial and academic sectors. Both human/machine and human/human communications can benefit from the application of SLU, using differing tasks and approaches to better understand and utilize such communications. This book covers the state-of-the-art approaches for the most popular SLU tasks with chapters written by well-known researchers in the respective fields. Key features include: Presents a fully integrated view of the two distinct disciplines of speech processing and language processing for SLU tasks. Defines what is possible today for SLU as an enabling technology for

enterprise (e.g., customer care centers or company meetings), and consumer (e.g., entertainment, mobile, car, robot, or smart environments) applications and outlines the key research areas. Provides a unique source of distilled information on methods for computer modeling of semantic information in human/machine and human/human conversations. This book can be successfully used for graduate courses in electronics engineering, computer science or computational linguistics. Moreover, technologists interested in processing spoken communications will find it a useful source of collated information of the topic drawn from the two distinct disciplines of speech processing and language processing under the new area of SLU.

Based on a NATO Advanced Study Institute held in 1993, this book addresses recent advances in automatic speech recognition and speech coding. The book contains contributions by many of the most outstanding researchers from the best laboratories worldwide in the field. The contributions have been grouped into five parts: on acoustic modeling; language modeling; speech processing,

analysis and synthesis; speech coding; and vector quantization and neural nets. For each of these topics, some of the best-known researchers were invited to give a lecture. In addition to these lectures, the topics were complemented with discussions and presentations of the work of those attending. Altogether, the reader is given a wide perspective on recent advances in the field and will be able to see the trends for future work.

The iPhone may be the world's coolest computer, but it's still a computer, with all of the complexities. iPhone: The Missing Manual is a illustrated guide to the tips, shortcuts, and workarounds that will turn you, too, into an iPhone master. This updated guide shows you everything you need to know about the new features and user interface of iOS 9 for the iPhone. This easy-to-use book will help you accomplish everything from web browsing to watching videos so you can get the most out of your iPhone.

Intelligent Speech Signal Processing investigates the utilization of speech analytics across several systems and real-world activities, including sharing data analytics, creating

collaboration networks between several participants, and implementing video-conferencing in different application areas. Chapters focus on the latest applications of speech data analysis and management tools across different recording systems. The book emphasizes the multidisciplinary nature of the field, presenting different applications and challenges with extensive studies on the design, development and management of intelligent systems, neural networks and related machine learning techniques for speech signal processing. Highlights different data analytics techniques in speech signal processing, including machine learning and data mining. Illustrates different applications and challenges across the design, implementation and management of intelligent systems and neural networks techniques for speech signal processing. Includes coverage of bi-modal speech recognition, voice activity detection, spoken language and speech disorder identification, automatic speech to speech summarization, and convolutional neural networks.

Automatic speech recognition (ASR) systems are

finding increasing use in everyday life. Many of the commonplace environments where the systems are used are noisy, for example users calling up a voice search system from a busy cafeteria or a street. This can result in degraded speech recordings and adversely affect the performance of speech recognition systems. As the use of ASR systems increases, knowledge of the state-of-the-art in techniques to deal with such problems becomes critical to system and application engineers and researchers who work with or on ASR technologies. This book presents a comprehensive survey of the state-of-the-art in techniques used to improve the robustness of speech recognition systems to these degrading external influences. Key features: Reviews all the main noise robust ASR approaches, including signal separation, voice activity detection, robust feature extraction, model compensation and adaptation, missing data techniques and recognition of reverberant speech. Acts as a timely exposition of the topic in light of more widespread use in the future of ASR technology in challenging environments. Addresses robustness issues and sig-

nal degradation which are both key requirements for practitioners of ASR. Includes contributions from top ASR researchers from leading research units in the field.

Research in the field of automatic speech and speaker recognition has made a number of significant advances in the last two decades, influenced by advances in signal processing, algorithms, architectures, and hardware. These advances include: the adoption of a statistical pattern recognition paradigm; the use of the hidden Markov modeling framework to characterize both the spectral and the temporal variations in the speech signal; the use of a large set of speech utterance examples from a large population of speakers to train the hidden Markov models of some fundamental speech units; the organization of speech and language knowledge sources into a structural finite state network; and the use of dynamic programming based heuristic search methods to find the best word sequence in the lexical network corresponding to the spoken utterance. Automatic Speech and Speaker Recognition: Advanced Topics groups together in a single volume

a number of important topics on speech and speaker recognition, topics which are of fundamental importance, but not yet covered in detail in existing textbooks. Although no explicit partition is given, the book is divided into five parts: Chapters 1-2 are devoted to technology overviews; Chapters 3-12 discuss acoustic modeling of fundamental speech units and lexical modeling of words and pronunciations; Chapters 13-15 address the issues related to flexibility and robustness; Chapter 16-18 concern the theoretical and practical issues of search; Chapters 19-20 give two examples of algorithm and implementational aspects for recognition system realization. Audience: A reference book for speech researchers and graduate students interested in pursuing potential research on the topic. May also be used as a text for advanced courses on the subject.

Automatic Speech Recognition (ASR) is undoubtedly one of the most important and interesting applications in the cutting-edge era of Deep-learning deployment, especially in the mobile segment. Fast and accurate ASR comes at a high energy cost, requiring huge

memory storage and computational power, which is not affordable for the tiny power budget of mobile devices. Hardware acceleration can reduce power consumption of ASR systems as well as reducing its memory pressure, while delivering high-performance. In this thesis, we present a customized accelerator for large-vocabulary, speaker-independent, continuous speech recognition. A state-of-the-art ASR system consists of two major components: acoustic-scoring using DNN and speech-graph decoding using Viterbi search. As the first step, we focus on the Viterbi search algorithm, that represents the main bottleneck in the ASR system. The accelerator includes some innovative techniques to improve the memory subsystem, which is the main bottleneck for performance and power, such as a prefetching scheme and a novel bandwidth saving technique tailored to the needs of ASR. Furthermore, as the speech graph is vast taking more than 1-Gigabyte memory space, we propose to change its representation by partitioning it into several sub-graphs and perform an on-the-fly composition during the Viterbi

run-time. This approach together with some simple yet efficient compression techniques result in 31x memory footprint reduction, providing 155x real-time speedup and orders of magnitude power and energy saving compared to CPUs and GPUs. In the next step, we propose a novel hardware-based ASR system that effectively integrates a DNN accelerator for the pruned/quantized models with the Viterbi accelerator. We show that, when either pruning or quantizing the DNN model used for acoustic scoring, ASR accuracy is maintained but the execution time of the ASR system is increased by 33%. Although pruning and quantization improves the efficiency of the DNN, they result in a huge increase of activity in the Viterbi search since the output scores of the pruned model are less reliable. In order to avoid the aforementioned increase in Viterbi search workload, our system loosely selects the N-best hypotheses at every time step, exploring only the N most likely paths. Our final solution manages to efficiently combine both DNN and Viterbi accelerators using all their optimizations, delivering 222x real-time ASR with a small power

budget of 1.26 Watt, small memory footprint of 41 MB, and a peak memory bandwidth of 381 MB/s, being amenable for low-power mobile platforms.