Speech Processing Rabiner Solution

Brief Notes in Advanced DSPGeophysical Signal AnalysisMachine Learning for Audio, Image and Video AnalysisSpringer Handbook of Speech ProcessingLinear Predictive Coding and the Internet ProtocolStatistical Methods for Speech RecognitionProgress in Nonlinear Speech Processing2003 IEEE Workshop on Applications of Signal Processing to Audio and AcousticsFast Algorithms for Signal ProcessingSpeech Coding AlgorithmsCircuits, Signals, and Speech and Image ProcessingFundamentals of Speech RecognitionPractical Digital Signal ProcessingDigital Processing of Speech SignalsBiometric SolutionsProgress in Pattern Recognition, Image Analysis, Computer Vision, and ApplicationsDigital Speech Processing Using MatlabAdvanced Signal Processing and Digital Noise ReductionIntroduction to Digital Speech ProcessingSignal Processing for CommunicationsDiscrete-Time Signal ProcessingMultidimensional Digital Signal ProcessingIntroducing Speech and Language ProcessingConnectionist Speech RecognitionWireless CommunicationsThe Digital Signal Processing HandbookTheory and Application of Digital Signal ProcessingVoice Communication Between Humans and MachinesApplied Signal ProcessingSignal AnalysisHidden Markov Models for BioinformaticsSpeech and Language Processing for Human-Machine CommunicationsSupplement: Introduction to Signal Processing & Computer Based Exercise Signal Processing Using MATLAB Version 5 Pkg. - IntroductiMobile Speech and Advanced Natural Language SolutionsSpoken Language ProcessingSelected Papers in Digital Signal Processing, IITheory and Applications of Digital Speech ProcessingMultirate Filtering for Digital Signal Processing: MATLAB ApplicationsAdvanced Algorithms and Architectures for Signal Processing IISpeech & Language Processing

Brief Notes in Advanced DSP

Remarkable progress is being made in spoken language processing, but many powerful techniques have remained hidden in conference proceedings and academic papers, inaccessible to most practitioners. In this book, the leaders of the Speech Technology Group at Microsoft Research share these advances -- presenting not just the latest theory, but practical techniques for building commercially viable products.KEY TOPICS: Spoken Language Processing draws upon the latest advances and techniques from multiple fields: acoustics, phonology, phonetics, linguistics, semantics, pragmatics, computer science, electrical engineering, mathematics, syntax, psychology, and beyond. The book begins by presenting essential background on speech production and perception, probability and information theory, and pattern recognition. The authors demonstrate how to extract useful information from the speech signal; then present a variety of contemporary speech recognition techniques, including hidden Markov models, acoustic and language modeling, and techniques for improving resistance to environmental noise. Coverage includes decoders, search algorithms, large vocabulary speech recognition techniques, text-to-speech, spoken language dialog management, user interfaces, and interaction with non-speech interface modalities. The authors also present detailed case studies based on Microsoft's advanced prototypes, including the Whisper speech recognizer, Whistler text-to-speech system, and MiPad handheld computer.MARKET: For anyone involved with planning, designing, building, or purchasing spoken language technology.

Geophysical Signal Analysis

Provides a clearly-written, concise and accessible introduction to speech and language processing, with accompanying software.

Machine Learning for Audio, Image and Video Analysis

Speech coding is a highly mature branch of signal processing deployed in products such as cellular phones, communication devices, and more recently, voice over internet protocol This book collects many of the techniques used in speech coding and presents them in an accessible fashion Emphasizes the foundation and evolution of standardized speech coders, covering standards from 1984 to the present The theory behind the applications is thoroughly analyzed and proved

Springer Handbook of Speech Processing

This text, an introduction to geophysical signal analysis, is concerned with the construction, analysis, and interpretation of mathematical and statistical models. In general, it is intended to provide material of interest to upper undergraduate students in mathematics, science, and engineering. Much of this book requires only a knowledge of elementary algebra. However, at some points, a familiarity with elementary calculus and matrix algebra is needed. The practical use of the concepts and techniques developed is illustrated by numerous applications. Care has been taken to choose examples that are of interest to a variety of readers. Therefore, the book contains material of interest to both geophysicists and those engaged in digital signal analysis in disciplines other than geophysics. This book is a reprint of the 1980 Prentice-Hall volume of the same title.

Linear Predictive Coding and the Internet Protocol

Applied Signal Processing: A MATLAB-Based Proof of Concept benefits readers by including the teaching background of experts in various applied signal processing fields and presenting them in a project-oriented framework. Unlike many other MATLAB-based textbooks which only use MATLAB to illustrate theoretical aspects, this book provides fully commented MATLAB code for working proofs-of-concept. The MATLAB code provided on the accompanying online files is the very heart of the material. In addition each chapter offers a functional introduction to the theory required to understand the code as well as a formatted presentation of the contents and outputs of the MATLAB code. Each chapter exposes how digital signal processing is applied for solving a real engineering problem used in a consumer product. The chapters are organized with a description of the problem in its applicative context and a functional review of the theory related to its solution appearing first. Equations are only used for a precise description of the problem and its final solutions. Then a step-by-step MATLAB-based proof of concept, with full code, graphs, and comments follows. The solutions are simple enough for readers with general signal processing background to understand and they use state-of-the-art signal processing principles. Applied Signal Processing: A MATLAB-Based Proof of Concept is an ideal companion for most signal processing course books. It can be used for preparing student labs and projects.

Statistical Methods for Speech Recognition

Bayesian Estimation and classification. Hidden markov models. Wiener filters. Kalman and adaptive least squared error filters.

Progress in Nonlinear Speech Processing

The field of digital signal processing (DSP) has spurred developments from basic theory of discrete-time signals and processing tools to diverse applications in telecommunications, speech and acoustics, radar, and video. This volume provides an accessible reference, offering theoretical and practical information to the audience of DSP users. This immense compilation outlines both introductory and specialized aspects of information-bearing signals in digital form, creating a resource relevant to the expanding needs of the engineering community. It also explores the use of computers and special-purpose digital hardware in extracting information or transforming signals in advantageous ways. Impacted areas presented include: Telecommunications Computer engineering Acoustics Seismic data analysis DSP software and hardware Image and video processing Remote sensing Multimedia applications Medical technology Radar and sonar applications This authoritative collaboration, written by the foremost researchers and practitioners in their fields, comprehensively presents the range of DSP: from theory to application, from algorithms to hardware.

2003 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics

Offers a well-rounded, mathematical approach to problems in signal interpretation using the latest time, frequency, and mixed-domain methods Equally useful as a reference, an up-to-date review, a learning tool, and a resource for signal analysis techniques Provides a gradual introduction to the mathematics so that the less mathematically adept reader will not be overwhelmed with instant hard analysis Covers Hilbert spaces, complex analysis, distributions, random signals, analog Fourier transforms, and more

Fast Algorithms for Signal Processing

Theory and Applications of Digital Speech Processing is ideal for graduate students in digital signal processing, and undergraduate students in Electrical and Computer Engineering. With its clear, up-to-date, hands-on coverage of digital speech processing, this text is also suitable for practicing engineers in speech processing. This new text presents the basic concepts and theories of speech processing with clarity and currency, while providing hands-on computer-based laboratory experiences for students. The material is organized in a manner that builds a strong foundation of basics first, and then concentrates on a range of signal processing methods for representing and processing the speech signal.

Speech Coding Algorithms

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.

Circuits, Signals, and Speech and Image Processing

Science fiction has long been populated with conversational computers and robots. Now, speech synthesis and recognition have matured to where a wide range of real-world applications-from serving people with disabilities to boosting the nation's competitiveness-are within our grasp. Voice Communication Between Humans and Machines takes the first interdisciplinary look at what we know about voice processing, where our technologies stand, and what the future may hold for this fascinating field. The volume integrates theoretical, technical, and practical views from world-class experts at leading research centers around the world, reporting on the scientific bases behind human-machine voice communication, the state of the art in computerization, and progress in user friendliness. It offers an up-to-date treatment of technological progress in key areas: speech synthesis, speech recognition, and natural language understanding. The book also explores the emergence of the voice processing industry and specific opportunities in telecommunications and other businesses, in military and government operations, and in assistance for the disabled. It outlines, as well, practical issues and research questions that must be resolved if machines are to become fellow problem-solvers along with humans. Voice Communication Between Humans and Machines provides a comprehensive understanding of the field of voice processing for engineers, researchers, and business executives, as well as speech and hearing specialists, advocates for people with disabilities, faculty and students, and interested individuals.

Fundamentals of Speech Recognition

Practical Digital Signal Processing

This book constitutes of the major results of the EU COST (European Cooperation in the field of Scientific and Technical Research) Action 277: NSP - Nonlinear Speech Processing - running from April 2001 to June 2005. The results were presented at the last meeting of the management committee of COST Action 277, held in Heraklion, Crete, Greece on September 20-23, 2005 during the Workshop on Nonlinear Speech Processing, WNSP 2005. The 13 revised full papers in this state-of-the-art survey were carefully reviewed and selected for inclusion in the book and are preceded with an introductory leading-in by the editors. The articles present overviews of the four years research combining linear and non linear approaches for processing the speech signal. The aim of this book is to provide an additional and/or an alternative way to the traditional approach of linear speech processing and be mainly used by the researcher working in the domain. The papers cover areas such as speech analysis for speech synthesis, speech recognition, speech-non speech discrimination and voice quality assessment, speaker recognition/verification from a natural or modified speech signal, speech recognition, speech enhancement, and emotional state detection.

Digital Processing of Speech Signals

Biometric Solutions

"This book covers basic and the advanced approaches in the design and implementation of multirate filtering"--Provided by publisher.

Progress in Pattern Recognition, Image Analysis, Computer Vision, and Applications

With a novel, less classical approach to the subject, the authors have written a book with the conviction that signal processing should be taught to be fun. The treatment is therefore less focused on the mathematics and more on the conceptual aspects, the idea being to allow the readers to think about the subject at a higher conceptual level, thus building the foundations for more advanced topics. The book remains an engineering text, with the goal of helping students solve real-world problems. In this vein, the last chapter pulls together the individual topics as discussed throughout the book into an in-depth look at the development of an end-to-end communication system, namely, a modem for communicating digital information over an analog channel.

Digital Speech Processing Using Matlab

Advanced Signal Processing and Digital Noise Reduction

Digital Speech Processing Using Matlab deals with digital speech pattern recognition, speech production model, speech feature extraction, and speech compression. The book is written in a manner that is suitable for beginners pursuing basic research in digital speech processing. Matlab illustrations are provided for most topics to enable better understanding of concepts. This book also deals with the basic pattern recognition techniques (illustrated with speech signals using Matlab) such as PCA, LDA, ICA, SVM, HMM, GMM, BPN, and KSOM.

Introduction to Digital Speech Processing

In December 1974 the first realtime conversation on the ARPAnet took place between Culler- Harrison Incorporated in Goleta, California, and MIT Lincoln Laboratory in Lexington, Massachusetts. This was the first successful application of realtime digital speech communication over a packet network and an early milestone in the explosion of realtime signal processing of speech, audio, images, and video that we all take for granted today. It could be considered as the first voice over Internet Protocol (VoIP), except that the Internet Protocol (IP) had not yet been established. In fact, the interest in realtime signal processing had an indirect, but major, impact on the development of IP. This is the story of the development of linear predictive coded (LPC) speech and how it came to be used in the first successful packet speech experiments.

Several related stories are recounted as well. The history is preceded by a tutorial on linear prediction methods which incorporates a variety of views to provide context for the stories. This part is a technical survey of the fundamental ideas of linear prediction that are important for speech processing, but the development departs from traditional treatments and takes advantage of several shortcuts, simplifications, and unifications that come with years of hindsight. In particular, some of the key results are proved using short and simple techniques that are not as well known as they should be, and it also addresses some of the common assumptions made when modeling random signals. The reader interested only in the history and already familiar with or uninterested in the technical details of linear prediction and speech may skip Part I entirely.

Signal Processing for Communications

This volume comprises the select proceedings of the annual convention of the Computer Society of India. Divided into 10 topical volumes, the proceedings present papers on state-of-the-art research, surveys, and succinct reviews. The volumes cover diverse topics ranging from communications networks to big data analytics, and from system architecture to cyber security. This volume focuses on Speech and Language Processing for Human-Machine Communications. The contents of this book will be useful to researchers and students alike.

Discrete-Time Signal 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.

Multidimensional Digital Signal Processing

The purpose of this book is to give a thorough and systematic introduction to probabilistic modeling in bioinformatics. The book contains a mathematically strict and extensive presentation of the kind of probabilistic models that have turned out to be useful in genome analysis. Questions of parametric inference, selection between model families, and various architectures are treated. Several examples are given of known architectures (e.g., profile HMM) used in genome analysis. Audience: This book will be of interest to advanced undergraduate and graduate students with a fairly limited background in probability theory, but otherwise well trained in mathematics and already familiar with at least some of the techniques of algorithmic sequence analysis.

Introducing Speech and Language Processing

Biometric Solutions for Authentication in an E-World provides a collection of sixteen chapters containing tutorial articles and new material in a unified manner. This includes the basic concepts, theories, and characteristic features of integrating/formulating different facets of biometric solutions for authentication, with recent developments and significant applications in an E-world. This book provides the reader with a basic concept of biometrics, an in-depth discussion exploring biometric technologies in various applications in an E-world. It also includes a detailed description of typical biometric-based security systems and up-to-date coverage of how these issues are developed. Experts from all over the world demonstrate the various ways this integration can be made to efficiently design methodologies, algorithms, architectures, and implementations for biometric-based applications in an E-world.

Connectionist Speech Recognition

Wireless Communications

This text on signal processing covers topics such as active echo and noise control, electronic musics, aids for auditory impairment, internet audio and applications, array processing for tranducers, music information retrieval, audio coding, signal processing for performing arts medicine and more.

The Digital Signal Processing Handbook

In two editions spanning more than a decade, The Electrical Engineering Handbook stands as the definitive reference to the multidisciplinary field of electrical engineering. Our knowledge continues to grow, and so does the Handbook. For the third edition, it has expanded into a set of six books carefully focused on a specialized area or field of study. Each book represents a concise yet definitive collection of key concepts, models, and equations in its respective domain, thoughtfully gathered for convenient access. Circuits, Signals, and Speech and Image Processing presents all of the basic information related to electric circuits and components, analysis of circuits, the use of the Laplace transform, as well as signal, speech, and image processing using filters and algorithms. It also examines emerging areas such as text-to-speech synthesis, real-time processing, and embedded signal processing. Each article includes defining terms, references, and sources of further information. Encompassing the work of the world's foremost experts in their respective specialties, Circuits, Signals, and Speech and Image Processing features the latest developments, the broadest scope of coverage, and new material on biometrics.

Theory and Application of Digital Signal Processing

Voice Communication Between Humans and Machines

Introduction to Digital Speech Processing highlights the central role of DSP techniques in modern speech communication research and applications. It presents a comprehensive overview of digital speech processing that ranges from the basic nature of the speech signal, through a variety of methods of representing speech in digital form, to applications in voice communication and automatic synthesis and recognition of speech. Introduction to Digital Speech Processing provides the reader with a practical introduction to the wide range of important concepts that comprise the field of digital speech processing. It serves as an invaluable reference for students embarking on speech research as well as the experienced researcher already working in the field, who can utilize the book as a reference guide.

Applied Signal Processing

This second edition focuses on audio, image and video data, the three main types of input that machines deal with when interacting with the real world. A set of appendices provides the reader with self-contained introductions to the mathematical background necessary to read the book. Divided into three main parts, From Perception to Computation introduces methodologies aimed at representing the data in forms suitable for computer processing, especially when it comes to audio and images. Whilst the second part, Machine Learning includes an extensive overview of statistical techniques aimed at addressing three main problems, namely classification (automatically assigning a data sample to one of the classes belonging to a predefined set), clustering (automatically grouping data samples according to the similarity of their properties) and sequence analysis (automatically mapping a sequence of observations into a sequence of human-understandable symbols). The third part Applications shows how the abstract problems defined in the second part underlie technologies capable to perform complex tasks such as the recognition of hand gestures or the transcription of handwritten data. Machine Learning for Audio, Image and Video Analysis is suitable for students to acquire a solid background in machine learning as well as for practitioners to deepen their knowledge of the state-of-the-art. All application chapters are based on publicly available data and free software packages, thus allowing readers to replicate the experiments.

Signal Analysis

The aim of this book is to introduce the general area of Digital Signal Processing from a practical point of view with a working minimum of mathematics. The emphasis is placed on the practical applications of DSP: implementation issues, tricks and pitfalls. Intuitive explanations and appropriate examples are used to develop a fundamental understanding of DSP theory, laying a firm foundation for the reader to pursue the matter further. The reader will develop a clear understanding of DSP technology in a variety of fields from process control to communications. * Covers the use of DSP in different engineering sectors, from communications to process control * Ideal for a wide audience wanting to take advantage of the strong movement towards digital signal processing techniques in the engineering world * Includes numerous practical exercises and diagrams covering many of the fundamental aspects of digital signal processing

Hidden Markov Models for Bioinformatics

This book reflects decades of important research on the mathematical foundations of speech recognition. It focuses on underlying statistical techniques such as hidden Markov models, decision trees, the expectation-maximization algorithm, information theoretic goodness criteria, maximum entropy probability estimation, parameter and data clustering, and smoothing of probability distributions. The author's goal is to present these principles clearly in the simplest setting, to show the advantages of self-organization from real data, and to enable the reader to apply the techniques.

Speech and Language Processing for Human-Machine Communications

Based on the authors’ research in Fourier analysis, Brief Notes in Advanced DSP: Fourier Analysis with MATLAB® addresses many concepts and applications of digital signal processing (DSP). The included MATLAB® codes illustrate how to apply the ideas in practice. The book begins with the basic concept of the discrete Fourier transformation and its properties. It then describes lifting schemes, integer transformations, the discrete cosine transform, and the paired transform method for calculating the discrete Hadamard transform. The text also examines the decomposition of the 1D signal by so-called section basis signals as well as new forms of 2D signal/image representation and decomposition by direction signals/images. Focusing on Fourier transform wavelets and Givens–Haar transforms, the last chapter discusses the problem of signal multiresolution. This book presents numerous interesting problems and concepts of unitary transformations, such as the Fourier, Hadamard, Hartley, Haar, paired, cosine, and new signal-induced transformations. It aids readers in using new forms and methods of signals and images in the frequency and frequency-and-time domains.

Supplement: Introduction to Signal Processing & Computer Based Exercise Signal Processing Using MATLAB Version 5 Pkg. - Introducti

"Professor Andreas F. Molisch, renowned researcher and educator, has put together the comprehensive book, Wireless Communications. The second edition, which includes a wealth of new material on important topics, ensures the role of the text as the key resource for every student, researcher, and practitioner in the field." —Professor Moe Win, MIT, USA Wireless communications has grown rapidly over the past decade from a niche market into one of the most important, fast moving industries. Fully updated to incorporate the latest research and developments, Wireless Communications, Second Edition provides an authoritative overview of the principles and applications of mobile communication technology. The author provides an in-depth analysis of current treatment of the area, addressing both the traditional elements, such as Rayleigh fading, BER in flat fading channels, and equalisation, and more recently emerging topics such as multi-user detection in CDMA systems, MIMO systems, and cognitive radio. The dominant wireless standards; including cellular, cordless and wireless LANs; are discussed. Topics featured include: wireless propagation channels, transceivers and signal processing, multiple access and advanced transceiver schemes, and standardised wireless systems. Combines mathematical descriptions with intuitive explanations of the physical facts, enabling readers to acquire a deep understanding of the subject. Includes new chapters on cognitive radio, cooperative communications and relaying, video coding, 3GPP Long Term Evolution, and WiMax; plus significant new sections on multi-user MIMO, 802.11n, and information theory. Companion website featuring: supplementary material on 'DECT', solutions manual and presentation slides for instructors, appendices, list of abbreviations and other useful resources.

Mobile Speech and Advanced Natural Language Solutions

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.

Spoken Language Processing

Multidimensional signals and systems. Discrete fourier analysis of multidimensional signals. Design and implementation of two-dimensional fir filters. Multidimensional recursive systems. Design and implementation of two-dimensional iir filters. Processing signals carried by propagation waves. Inverse problems.

Selected Papers in Digital Signal Processing, II

Theory and Applications of Digital Speech Processing

Pattern recognition is a central topic in contemporary computer sciences, with continuously evolving topics, challenges, and methods, including machine learning, content-based image retrieval, and model- and knowledge-based - proaches, just to name a few. The Iberoamerican Congress on Pattern Recog- tion (CIARP) has become established as a high-quality conference, highlighting the recent evolution of the domain. These proceedings include all papers presented during the 15th edition of this conference, held in Sao Paulo, Brazil, in November 2010. As was the case for previous conferences, CIARP 2010 attracted parti- pants from around the world with the aim of promoting and disseminating - going research on mathematical methods and computing techniques for pattern recognition, computer vision, image analysis, and speech recognition, as well as their applications in such diverse areas as robotics, health, entertainment, space exploration, telecommunications, data mining, document analysis, and natural language processing and recognition, to name only a few of them. Moreover, it provided a forum for scienti?c research, experience exchange, sharing new kno- edge and increasing cooperation between research groups in pattern recognition and related areas. It is important to underline that these conferences have contributed sign- icantly to the growth of national associations for pattern recognition in the Iberoamerican region, all of them as members of the International Association for Pattern Recognition (IAPR).

Multirate Filtering for Digital Signal Processing: MATLAB Applications

Advanced Algorithms and Architectures for Signal Processing II

Efficient signal processing algorithms are important for embedded and power-limited applications since, by reducing the number of computations, power consumption can be reduced significantly. Similarly, efficient algorithms are also critical to very large scale applications such as video processing and four-dimensional medical imaging. This self-contained guide, the only one of its kind, enables engineers to find the optimum fast algorithm for a specific application. It presents a broad range of computationally-efficient algorithms, describes their structure and implementation, and compares their relative strengths for given problems. All the necessary background mathematics is included and theorems are rigorously proved, so all the information needed to learn and apply the techniques is provided in one convenient guide. With this practical reference, researchers and practitioners in electrical engineering, applied mathematics, and computer science can reduce power dissipation for low-end applications of signal processing, and extend the reach of high-end applications.

Speech & Language Processing

Provides a theoretically sound, technically accurate, and complete description of the basic knowledge and ideas that constitute a modern system for speech recognition by machine. Covers production, perception, and acoustic-phonetic characterization of the speech signal; signal processing and analysis methods for speech recognition; pattern comparison techniques; speech recognition system design and implementation; theory and implementation of hidden Markov models; speech recognition based on connected word models; large vocabulary continuous speech recognition; and task- oriented application of automatic speech recognition. For practicing engineers, scientists, linguists, and programmers interested in speech recognition.