Hidden Markov Models For Time Series

Introduction to Stochastic Processes with R Springer Science & Business Media
Time series data analysis is increasingly important due to the massive production of such data through the internet of things, the digitalization of healthcare, and the rise of smart cities. As continuous monitoring and data collection become more common, the need for competent time series analysis with both statistical and machine learning techniques will increase. Covering innovations in time series data analysis and use cases from the real world, this practical guide will help you solve the most common data engineering and analysis challengesin time series, using both traditional statistical and modern machine learning techniques. Author Aileen Nielsen offers an accessible, well-rounded introduction to time series in both R and Python that will have data scientists, software engineers, and researchers up and running quickly. You'll get the guidance you need to confidently: Find and wrangle time series data Undertake exploratory time series data analysis Store temporal data Simulate time series data Generate and select features for a time series Measure error Forecast and classify time series with machine or deep learning Evaluate accuracy and performance

Speech & Language Processing Cambridge University Press

In the current age of information technology, the issues of distributing and utilizing images efficiently and effectively are of substantial concern. Solutions to many of the problems arising from these issues are provided by techniques of image processing, among which segmentation and compression are topics of this book. Image segmentation is a process for dividing an image into its constituent parts. For block-based segmentation using statistical classification, an image is divided into blocks and a feature vector is formed for each block by grouping statistics of its pixel intensities. Conventional block-based segmentation algorithms classify each block separately, assuming independence of feature vectors. Image Segmentation and Compression Using Hidden Markov Models presents a new algorithm that models the statistical dependence among image blocks by two dimensional hidden Markov models (HMMs). Formulas for estimating the model according to the maximum likelihood criterion are derived from the EM algorithm. To segment an image, optimal classes are searched jointly for all the blocks by the maximum a posteriori (MAP) rule. The 2-D HMM is extended to multiresolution so that more context information is exploited in classification and fast progressive segmentation schemes can be formed naturally. The second issue addressed in the book is the design of joint compression and classification systems using the 2-D HMM and vector quantization. A classifier designed with the side goal of good compression often outperforms one aimed solely at classification because overfitting to training data is suppressed by vector quantization. Image Segmentation and Compression Using Hidden Markov Models is an essential reference source for researchers and engineers working in statistical signal processing or image processing, especially those who are interested in hidden Markov models. It is also of value to those working on statistical modeling.

Semi-Markov Chains and Hidden Semi-Markov Models toward Applications Springer Science & Business Media
"The thesis presents two Hidden Markov Model (HMM) based methodologies for the analysis and prediction of financial time series. Both methodologies are symbolic i.e., the time series is discretized into a sequence of symbols and future symbols are predicted based on the past symbols"--Abstract, leaf iii.

Efficient Learning Machines CRC Press
Presents algorithms for using HMMs and explains the derivation of those algorithms for the dynamical systems community.

Hidden Markov Models for Time Series CRC Press
Hidden Markov Models for Time Series: An Introduction Using R, Second Edition illustrates the great flexibility of hidden Markov models (HMMs) as general-purpose models for time series data. The book provides a broad understanding of the models and their uses. After presenting the basic model formulation, the book covers estimation, forecasting, decoding, prediction, model selection, and Bayesian inference for HMMs. Through examples and applications, the authors describe how to extend and generalize the basic model so that it can be applied in a rich variety of situations. The book demonstrates how HMMs can be applied to a wide range of types of time series: continuous-valued, circular, multivariate, binary, bounded and unbounded counts, and categorical observations. It also discusses how to employ the freely available computing environment R to carry out the computations. Features Presents an accessible overview of HMMs Explores a variety of applications in ecology, finance, epidemiology, climatology, and sociology Includes numerous theoretical and programming exercises Provides most of the analysed data sets online New to the second edition A total of five chapters on extensions, including HMMs for longitudinal data, hidden semi-Markov models and models with continuous-valued state process New case studies on animal movement, rainfall occurrence and capture recapture data "

Hidden Markov Models in Finance O'Reilly Media
An introduction to stochastic processes through the use of R Introduction to Stochastic Processes with R is an accessible and well-balanced presentation of the theory of stochastic processes, with an emphasis on real-world applications of probability theory in the natural and social sciences. The use of simulation, by means of the popular statistical freeware R, makes theoretical results come alive with practical, hands-on demonstrations. Written by a highly-qualified expert in the field, the author presents numerous examples from a wide array of disciplines, which are used to illustrate concepts and highlight computational and theoretical results. Developing readers' problem-solving skills and mathematical maturity, Introduction to Stochastic Processes with R features: Over 200 examples and 600 end-of-chapter exercises A tutorial for getting started with R, and appendices that contain review material in probability and matrix algebra Discussions of many timely and interesting supplemental topics including Markov chain Monte Carlo, random walk on graphs, card shuffling, Black-Scholes options pricing, applications in biology and genetics, cryptography, martingales, and stochastic calculus Introductions to mathematics as needed in order to suit readers at many mathematical levels A companion website that includes relevant data files as well as all R code and scripts used throughout the book Introduction to Stochastic Processes with R is an ideal textbook for an introductory course in stochastic processes.

The book is aimed at undergraduate and beginning graduate-level students in the science, technology, engineering, and mathematics disciplines. The book is also an excellent reference for applied mathematicians and statisticians who are interested in a review of the topic.
Machine Learning SIAM
A practical introduction perfect for final-year undergraduate and graduate students without a solid background in linear algebra and calculus.
Inference in Hidden Markov Models Morgan Kaufmann
Hidden Markov Models for Time Series: An Introduction Using R, Second Edition illustrates the great flexibility of hidden Markov models (HMMs) as general-purpose models for time series data. The book provides a broad understanding of the models and their uses. After presenting the basic model formulation, the book covers estimation, forecasting, decoding, prediction, model selection, and Bayesian inference for HMMs. Through examples and applications, the authors describe how to extend and generalize the basic model so that it can be applied in a rich variety of situations. The book demonstrates how HMMs can be applied to a wide range of types of time series: continuous-valued, circular, multivariate, binary, bounded and unbounded counts, and categorical observations. It also discusses how to employ the freely available computing environment R to carry out the computations. Features Presents an accessible overview of HMMs Explores a variety of applications in ecology, finance, epidemiology, climatology, and sociology Includes numerous theoretical and programming exercises Provides most of the analysed data sets online New to the second edition A total of five chapters on extensions, including HMMs for longitudinal data, hidden semi-Markov models and models with continuous-valued state process New case studies on animal movement, rainfall occurrence and capture-recapture data

Hidden Markov Models for Time Series Cambridge University Press
The study of animal movement has always been a key element in ecological science, because it is inherently linked to critical processes that scale from individuals to populations and communities to ecosystems. Rapid improvements in biotelemetry data collection and processing technology have given rise to a variety of statistical methods for characterizing animal movement. The book serves as a comprehensive reference for the types of statistical models used to study individual-based animal movement. Animal Movement is an essential reference for wildlife biologists, quantitative ecologists, and statisticians who seek a deeper understanding of modern animal movement models. A wide variety of modeling approaches are reconciled in the book using a consistent notation. Models are organized into groups based on how they treat the underlying spatio-temporal process of movement. Connections among approaches are highlighted to allow the reader to form a broader view of animal movement analysis and its associations with traditional spatial and temporal statistical modeling. After an initial overview examining the role that animal movement plays in ecology, a primer on spatial and temporal statistics provides a solid foundation for the remainder of the book. Each subsequent chapter outlines a fundamental type of statistical model utilized in the contemporary analysis of telemetry data for animal movement inference. Descriptions begin with basic traditional forms and sequentially build up to general classes of models in each category. Important background and technical details for each class of model are provided, including spatial point process models, discrete-time dynamic models, and continuous-time stochastic process models. The book also covers the essential elements for how to accommodate multiple sources of uncertainty, such as location error and latent behavior states. In addition to thorough descriptions of animal movement models, differences and connections among approaches are also emphasized to provide a broader perspective of approaches.

Latent Markov Models for Longitudinal Data Chapman and Hall/CRC
Machine learning techniques provide cost-effective alternatives to traditional methods for extracting underlying relationships between information and data and for predicting future events by processing existing information to train models. Efficient Learning Machines explores the major topics of machine learning, including knowledge discovery, classifications, genetic algorithms, neural networking, kernel methods, and biologically-inspired techniques. Mariette Awad and Rahul Khanna's synthetic approach weaves together the theoretical exposition, design principles, and practical applications of efficient machine learning. Their experiential emphasis, expressed in their close analysis of sample algorithms throughout the book, aims to equip engineers, students of engineering, and system designers to design and create new and more efficient machine learning systems. Readers of Efficient Learning Machines will learn how to recognize and analyze the problems that machine learning technology can solve for them, how to implement and deploy standard solutions to sample problems, and how to design new systems and solutions. Advances in computing performance, storage, memory, unstructured information retrieval, and cloud computing have coevolved with a new generation of machine learning paradigms and big data analytics, which the authors present in the conceptual context of their traditional precursors. Awad and Khanna explore current developments in the deep learning techniques of deep neural networks, hierarchical temporal memory, and cortical algorithms. Nature suggests sophisticated learning techniques that deploy simple rules to generate highly intelligent and organized behaviors with adaptive, evolutionary, and distributed properties. The authors examine the most popular biologically-inspired algorithms, together with a sample application to distributed datacenter management. They also discuss machine learning techniques for addressing problems of multi-objective optimization in which solutions in real-world systems are constrained and evaluated based on how well they perform with respect to multiple objectives in aggregate. Two chapters on support vector machines and their extensions focus on recent improvements to the

classification and regression techniques at the core of machine learning.

Theories, Concepts, and Applications for Engineers and System Designers Steven Taylor
Here is a work that adds much to the sum of our knowledge in a key area of science today. It is concerned with the estimation of discrete-time semi-Markov and hidden semi-Markov processes. A unique feature of the book is the use of discrete time, especially useful in some specific applications where the time scale is intrinsically discrete. The models presented in the book are specifically adapted to reliability studies and DNA analysis. The book is mainly intended for applied probabilists and statisticians interested in semi-Markov chains theory, reliability and DNA analysis, and for theoretical oriented reliability and bioinformatics engineers.

Symbolic Time Series Analysis Using Hidden Markov Models Springer
This book is a comprehensive treatment of inference for hidden Markov models, including both algorithms and statistical theory. Topics range from filtering and smoothing of the hidden Markov chain to parameter estimation, Bayesian methods and estimation of the number of states. In a unified way the book covers both models with finite state spaces and models with continuous state spaces (also called state-space models) requiring approximate simulation-based algorithms that are also described in detail. Many examples illustrate the algorithms and theory. This book builds on recent developments to present a self-contained view.

Practical Time Series Analysis Springer Science & Business Media
Drawing on the authors' extensive research in the analysis of categorical longitudinal data, Latent Markov Models for Longitudinal Data focuses on the formulation of latent Markov models and the practical use of these models. Numerous examples illustrate how latent Markov models are used in economics, education, sociology, and other fields. The R and MATLAB® routines used for the examples are available on the authors' website. The book provides you with the essential background on latent variable models, particularly the latent class model. It discusses how the Markov chain model and the latent class model represent a useful paradigm for latent Markov models. The authors illustrate the assumptions of the basic version of the latent Markov model and introduce maximum likelihood estimation through the Expectation-Maximization algorithm. They also cover constrained versions of the basic latent Markov model, describe the inclusion of the individual covariates, and address the random effects and multilevel extensions of the model. After covering advanced topics, the book concludes with a discussion on Bayesian inference as an alternative to maximum likelihood inference. As longitudinal data become increasingly relevant in many fields, researchers must rely on specific statistical and econometric models tailored to their application. A complete overview of latent Markov models, this book demonstrates how to use the models in three types of analysis: transition analysis with measurement errors, analyses that consider unobserved heterogeneity, and finding clusters of units and studying the transition between the clusters.

Bayesian Reasoning and Machine Learning Cambridge University Press
Discrete-valued time series are common in practice, but methods for their analysis are not well-known. In recent years, methods have been developed which are specifically designed for the analysis of discrete-valued time series. Hidden Markov and Other Models for Discrete-Valued Time Series introduces a new, versatile, and computationally tractable class of models, the "hidden Markov" models. It presents a detailed account of these models, then applies them to data from a wide range of diverse subject areas, including medicine, climatology, and geophysics. This book will be invaluable to researchers and postgraduate and senior undergraduate students in statistics. Researchers and applied statisticians who analyze time series data in medicine, animal behavior, hydrology, and sociology will also find this information useful.

Hidden Markov Models for Speech Recognition Springer
The Application of Hidden Markov Models in Speech Recognition presents the core architecture of a HMM-based LVCSR system and proceeds to describe the various refinements which are needed to achieve state-of-the-art performance.

Hidden Markov Models and Dynamical Systems CRC Press
Gain the confidence you need to apply machine learning in your daily work. With this practical guide, author Matthew Kirk shows you how to integrate and test machine learning algorithms in your code, without the academic subtext. Featuring graphs and highlighted code examples throughout, the book features tests with Python's Numpy, Pandas, Scikit-Learn, and SciPy data science libraries. If you're a software engineer or business analyst interested in data science, this book will help you: Reference real-world examples to test each algorithm through engaging, hands-on exercises Apply test-driven development (TDD) to write and run tests before you start coding Explore techniques for improving your machine-learning models with data extraction and feature development Watch out for the risks of machine learning, such as underfitting or overfitting data Work with K-Nearest Neighbors, neural networks, clustering, and other algorithms

Machine Learning: ECML 2007 CRC Press
Unleash the power of unsupervised machine learning in Hidden Markov Models using TensorFlow, pgmpy, and hmmlearn Key Features Build a variety of Hidden Markov Models (HMM) Create and apply models to any sequence of data to analyze, predict, and extract valuable insights Use natural language processing (NLP) techniques and 2D-HMM model for image segmentation Book Description Hidden Markov Model (HMM) is a statistical model based on the Markov chain concept. Hands-On Markov Models with Python helps you get to grips with HMMs and different inference algorithms by working on real-world problems. The hands-on examples explored in the book help you simplify the process flow in machine learning by using Markov model concepts, thereby making it accessible to everyone. Once you've covered the basic concepts of Markov chains, you'll get insights into Markov processes, models, and types with the help of practical examples. After grasping these fundamentals, you'll move on to learning about the different algorithms used in inferences and applying them in state and parameter inference. In addition to this, you'll explore the Bayesian approach of inference and learn how to apply it in HMMs. In further chapters, you'll discover how to use HMMs in time series analysis and natural language processing (NLP) using Python. You'll also learn to apply HMM to image processing using 2D-HMM to segment images. Finally, you'll understand how to apply HMM for reinforcement learning (RL) with the help of Q-Learning, and use this technique for single-stock and multi-stock algorithmic trading. By the end of this book, you will have grasped how to build your own

Markov and hidden Markov models on complex datasets in order to apply them to projects. What you will learn Explore a balance of both theoretical and practical aspects of HMM Implement HMMs using different datasets in Python using different packages Understand multiple inference algorithms and how to select the right algorithm to resolve your problems Develop a Bayesian approach to inference in HMMs Implement HMMs in finance, natural language processing (NLP), and image processing Determine the most likely sequence of hidden states in an HMM using the Viterbi algorithm Who this book is for Hands-On Markov Models with Python is for you if you are a data analyst, data scientist, or machine learning developer and want to enhance your machine learning knowledge and skills. This book will also help you build your own hidden Markov models by applying them to any sequence of data. Basic knowledge of machine learning and the Python programming language is expected to get the most out of the book

Hidden Semi-Markov Models CRC Press
The first unified treatment of time series modelling techniques spanning machine learning, statistics, engineering and computer science.

An Introduction Using R, Second Edition Springer Nature
A comprehensive introduction to machine learning that uses probabilistic models and inference as a unifying approach. Today's Web-enabled deluge of electronic data calls for automated methods of data analysis. Machine learning provides these, developing methods that can automatically detect patterns in data and then use the uncovered patterns to predict future data. This textbook offers a comprehensive and self-contained introduction to the field of machine learning, based on a unified, probabilistic approach. The coverage combines breadth and depth, offering necessary background material on such topics as probability, optimization, and linear algebra as well as discussion of recent developments in the field, including conditional random fields, L1 regularization, and deep learning. The book is written in an informal, accessible style, complete with pseudo-code for the most important algorithms. All topics are copiously illustrated with color images and worked examples drawn from such application domains as biology, text processing, computer vision, and robotics. Rather than providing a cookbook of different heuristic methods, the book stresses a principled model-based approach, often using the language of graphical models to specify models in a concise and intuitive way. Almost all the models described have been implemented in a MATLAB software package—PMTK (probabilistic modeling toolkit)—that is freely available online. The book is suitable for upper-level undergraduates with an introductory-level college math background and beginning graduate students.

Hidden Markov Models for Time Series Springer
Bayesian nonparametrics works - theoretically, computationally. The theory provides highly flexible models whose complexity grows appropriately with the amount of data. Computational issues, though challenging, are no longer intractable. All that is needed is an entry point: this intelligent book is the perfect guide to what can seem a forbidding landscape. Tutorial chapters by Ghosal, Lijoi and Prünster, Teh and Jordan, and Dunson advance from theory, to basic models and hierarchical modeling, to applications and implementation, particularly in computer science and biostatistics. These are complemented by companion chapters by the editors and Griffin and Quintana, providing additional models, examining computational issues, identifying future growth areas, and giving links to related topics. This coherent text gives ready access both to underlying principles and to state-of-the-art practice. Specific examples are drawn from information retrieval, NLP, machine vision, computational biology, biostatistics, and bioinformatics.