BAYESIAN TIME SERIES MODELS

'What's going to happen next?' Time series data hold the answers, and Bayesian methods represent the cutting edge in learning what they have to say. This ambitious book is the first unified treatment of the emerging knowledge-base in Bayesian time series techniques. Exploiting the unifying framework of probabilistic graphical models, the book covers approximation schemes, both Monte Carlo and deterministic, and introduces switching, multi-object, nonparametric and agent-based models in a variety of application environments. It demonstrates that the basic framework supports the rapid creation of models tailored to specific applications and gives insight into the computational complexity of their implementation.

The authors span traditional disciplines such as statistics and engineering and the more recently established areas of machine learning and pattern recognition. Readers with a basic understanding of applied probability, but no experience with time series analysis, are guided from fundamental concepts to the state of the art in research and practice.

BAYESIAN TIME SERIES MODELS

Edited by

David Barber
University College London

A. Taylan Cemgil
Boğaziçi University, Bebek, Istanbul

Silvia Chiappa
University of Cambridge

CAMBRIDGE
UNIVERSITY PRESS

CAMBRIDGE UNIVERSITY PRESS
Cambridge, New York, Melbourne, Madrid, Cape Town,
Singapore, São Paulo, Delhi, Tokyo, Mexico City

Cambridge University Press
The Edinburgh Building, Cambridge CB2 8RU, UK

Published in the United States of America by Cambridge University Press, New York

www.cambridge.org
Information on this title: www.cambridge.org/9780521196765

First published 2011

Printed in the United Kingdom at the University Press, Cambridge

A catalogue record for this publication is available from the British Library

Library of Congress Cataloguing in Publication data
Bayesian time series models / edited by David Barber, A. Taylan Cemgil, Silvia Chiappa.
p. cm.
ISBN 978-0-521-19676-5 (hardback)
1. Time-series analysis. 2. Bayesian statistical decision theory. I. Barber, David, 1968 Nov. 9–
II. Cemgil, Ali Taylan. III. Chiappa, Silvia. IV. Title.
QA280.B39 2011
519.5′5–dc22
2011008051

ISBN 978-0-521-19676-5 Hardback

Contents

Contributors

Cédric Archambeau, *Department of Computer Science, University College London*

Yves Atchadé *Department of Statistics, University of Michigan*

David Barber *Department of Computer Science, University College London*

A. Taylan Cemgil *Department of Computer Engineering Boğaziçi University*

Silvia Chiappa *Statistical Laboratory, University of Cambridge*

Idris A. Eckley *Department of Mathematics and Statistics, Lancaster University*

Paul Fearnhead *Department of Mathematics and Statistics, Lancaster University*

Gersende Fort *LTCI, CNRS – Telecom ParisTech*

Jurgen Van Gael *Department of Engineering, University of Cambridge*

Zoubin Ghahramani *Department of Engineering, University of Cambridge*

Simon J. Godsill *Signal Processing Laboratory, Department of Engineering, University of Cambridge*

Stefan Harmeling *MPI for Biological Cybernetics, Department Schölkopf, Tübingen*

Tom Heskes *Institute for Computing and Information Sciences, Radboud University Nijmegen*

Simon Hill *Signal Processing Laboratory, Department of Engineering, University of Cambridge*

Nick R. Jennings *School of Electronics and Computer Science, University of Southampton*

Adam M. Johansen *Department of Statistics, University of Warwick*

Hilbert J. Kappen *Donders' Institute for Neuroscience, Radboud University, Nijmegen*

Rebecca Killick *Department of Mathematics and Statistics, Lancaster University*

Risi Kondor *Center for the Mathematics of Information, California Institute of Technology*

Neil D. Lawrence *School of Computer Science, University of Manchester*

Jack Li *Signal Processing Laboratory, Department of Engineering, University of Cambridge*

Eric Moulines *LTCI, CNRS – Telecom ParisTech*

Manfred Opper *Technische Universität Berlin, Fakultät IV – Elektrotechnik und Informatik*

Michael A. Osborne *Department of Engineering Science, University of Oxford*

Sze Kim Pang *Signal Processing Laboratory, Department of Engineering, University of Cambridge*

Omiros Papaspiliopoulos *Department of Economics, Universitat Pompeu Fabra, Barcelona*

Pierre Priouret *Laboratoire de Probabilités et Modèles Aléatoires, Université P. & M. Curie, Paris*

John A. Quinn *Department of Computer Science, Makerere University*

Sarvapali D. Ramchurn *School of Electronics and Computer Science, University of Southampton*

Magnus Rattray *School of Computer Science, University of Manchester*

Stephen J. Roberts *Department of Engineering Science, University of Oxford*

Alex Rogers *School of Electronics and Computer Science, University of Southampton*

Maneesh Sahani *Gatsby Computational Neuroscience Unit, London*

François Septier *Institut TELECOM/ TELECOM Lille 1, France*

Sumeetpal S. Singh *Signal Processing Laboratory, Department of Engineering, University of Cambridge*

Amos Storkey *Institute for Adaptive and Neural Computation, University of Edinburgh*

Michalis K. Titsias *School of Computer Science, University of Manchester*

Marc Toussaint *Technische Universität Berlin*

Richard Eric Turner *Gatsby Computational Neuroscience Unit, London*

Nick Whiteley *Statistics Group, Department of Mathematics, University of Bristol*

Christopher K. I. Williams *Institute for Adaptive and Neural Computation, University of Edinburgh*

Onno Zoeter *Xerox Research Centre Europe, Meylan*

Preface

Probabilistic time series modelling

Time series are studied in a variety of disciplines and appear in many modern applications such as financial time series prediction, video-tracking, music analysis, control and genetic sequence analysis. This widespread interest at times obscures the commonalities in the developed models and techniques. A central aim of this book is to attempt to make modern time series techniques, specifically those based on probabilistic modelling, accessible to a broad range of researchers.

In order to achieve this goal, leading researchers that span the more traditional disciplines of statistics, control theory, engineering and signal processing, and the more recent areas of machine learning and pattern recognition, have been brought together to discuss advancements and developments in their respective fields. In addition, the book makes extensive use of the graphical models framework. This framework facilitates the representation of many classical models and provides insight into the computational complexity of their implementation. Furthermore, it enables to easily envisage new models tailored for a particular environment. For example, the book discusses novel state space models and their application in signal processing including condition monitoring and tracking. The book also describes modern developments in the machine learning community applied to more traditional areas of control theory.

The effective application of probabilistic models in the real world is gaining pace, largely through increased computational power which brings more general models into consideration through carefully developed implementations. As such, developing new models and associated approximate inference schemes is likely to remain an active area of research, with graphical models playing an important role in facilitating communication and guiding intuition. The book extensively discusses novel developments in approximate inference, including both deterministic and stochastic approximations.

The structure of the book

Chapter 1 gives a general introduction to probabilistic time series and explains how graphical models can be used to compactly represent classical models, such as the linear dynamical system and hidden Markov model. The chapter also discusses stochastic approximation schemes such as Markov chains and sequential Monte Carlo (particle filtering), and less well known deterministic approximation schemes such as variational methods.

The subsequent chapters are organised into six thematic parts: the first two deal with more theoretical issues related to approximate inference, while the remaining four deal with the development and application of novel models. More specifically, the parts are organised as follows.

Monte Carlo Monte Carlo methods are important and widespread techniques for approximate inference in probabilistic models. Chapter 2 gives a comprehensive introduction to adaptive Markov chain Monte Carlo methods. In time series models, a particularly relevant issue is that data often arrives sequentially, for which sequential Monte Carlo methods are appropriate. Chapter 3 gives a survey of recent developments in particle filtering. Chapter 4 presents the application of Monte Carlo methods to diffusion processes.

Deterministic approximations Chapter 5 discusses some characteristics of variational approximations, highlighting important aspects and difficulties idiosyncratic to time series models. Chapter 6 presents a novel deterministic approximate inference method for continuous-time Markov processes. Chapters 7 and 8 deal with inference in the important but computationally difficult switching linear dynamical system, and introduce specific deterministic inference schemes as improvements on classical methods.

Switching models Switching models assume that an underlying process may change from one parameter regime to another over time and may be used to model changes in the environment. In Chapter 9 switching models are applied to condition monitoring, in particular physiological monitoring. Chapter 10 reviews changepoint models, which are restricted switching models used to detect abrupt changes in time series, and gives insights into new applications.

Multi-object models A particularly active research area is the tracking of moving objects, such as crowds. In Chapter 11, a detailed discussion of a tracking model and associated stochastic inference method is given. In a similar vein, in Chapter 12 a framework for sequentially inferring how groups dynamically evolve is presented and applied to tracking the group behaviour of financial stocks. In Chapter 13 a different theoretical approach to multi-object tracking based on non-commutative harmonic analysis is discussed.

Nonparametric models In recent years flexible nonparametric models have been a particular focus of machine learning research. In this part, their extension to time series analysis is presented. Chapter 14 discusses sampling algorithms for Gaussian processes. Chapter 15 presents recent developments in nonparametric hidden Markov models, along with associated inference algorithms and applications. Chapter 16 introduces an application of a nonparametric time series model to multi-sensor time series prediction.

Agent-based models A recent viewpoint is to treat a control problem as an inference problem in an associated probabilistic model. This viewpoint is complementary to classical control theory and reinforcement learning and makes use of concepts familiar to probabilistic modellers, facilitating an entrance into this field. In Chapter 17 optimal control theory and the linear Bellman equation are discussed in relation to inference in a probabilistic model. In Chapter 18 the standard methods of learning in probabilistic models are applied to learning control policies in Markov decision problems.

Whom this book is for

This book will appeal to statisticians interested in modern aspects of time series analysis in areas bordering with engineering, signal processing and machine learning and how time series analysis is approached in those communities. For engineers and machine learners, the book has a wealth of insights into statistical approaches, particularly dealing with sampling techniques applied to difficult time series models. To follow the book, no specific knowledge of time series is required. However, readers are assumed to have a basic understanding of applied probability.

Acknowledgements

We are particularly grateful to the following people for their advice and comments on the book: Julien Cornebise, Mark Girolami, Andrew Golightly, Jim Griffin, Matt Hoffmanm, Antti Honkela, Jonathan Huang, Ajay Jasra, Jens Kober, Jan Peters, Sonia Petrone, Alan Qi, George Sermaidis, Olivier Stegle, Matt Taddy, Evangelos Theodorou, Yener Ulker.