

The background of the top half of the cover is a dark blue grid. Overlaid on the grid are several financial data visualizations: a candlestick chart at the top with white and grey bars and numerical price labels (14288, 14254, 15214, 13678, 13191, 13462, 12892); a smooth, wavy line graph in the middle; and a bar chart at the bottom. The text is positioned in the upper right corner.

EDITED BY

**ANTHONY  
O'HAGAN**


**MIKE  
WEST**

≡ The Oxford Handbook of  
**APPLIED BAYESIAN  
ANALYSIS**

# The Oxford Handbook of Applied Bayesian Analysis

Edited by

Anthony O'Hagan  
Mike West



OXFORD  
UNIVERSITY PRESS

**OXFORD**  
UNIVERSITY PRESS

Great Clarendon Street, Oxford, OX2 6DP,  
United Kingdom

Oxford University Press is a department of the University of Oxford.  
It furthers the University's objective of excellence in research, scholarship,  
and education by publishing worldwide. Oxford is a registered trade mark of  
Oxford University Press in the UK and in certain other countries

© Oxford University Press 2010

The moral rights of the authors have been asserted

First published 2010

First published in paperback 2013

Impression: 1

All rights reserved. No part of this publication may be reproduced, stored in  
a retrieval system, or transmitted, in any form or by any means, without the  
prior permission in writing of Oxford University Press, or as expressly permitted  
by law, by licence or under terms agreed with the appropriate reprographics  
rights organization. Enquiries concerning reproduction outside the scope of the  
above should be sent to the Rights Department, Oxford University Press, at the  
address above

You must not circulate this work in any other form  
and you must impose this same condition on any acquirer

Published in the United States of America by Oxford University Press  
198 Madison Avenue, New York, NY 10016, United States of America

British Library Cataloguing in Publication Data

Data available

Library of Congress Cataloging in Publication Data

Data available

ISBN 978-0-19-954890-3 (Hbk.)

ISBN 978-0-19-870317-4 (Pbk.)

Printed in Great Britain on acid-free paper by  
CPI Antony Rowe, Chippenham, Wiltshire

# The Oxford Handbook of Applied Bayesian Analysis

# Preface

## A Bayesian 21st century

The diversity of applications of modern Bayesian analysis at the start of the 21st century is simply enormous. From basic biology to frontier information technology, the applications of highly structured stochastic models of increasing realism – often with high-dimensional parameters and latent variables, multiple layers of hierarchically structured random effects, and nonparametric components – are increasingly routine. Much of the impetus behind this growth and success of applied Bayesian methods over the last 20 years has come from access to the increasingly rich array of advanced computational strategies for Bayesian analysis; this has led to increasing adoption of Bayesian methods from heavily practical and pragmatic perspectives.

Coupled with this evolution in the nature of applied statistical work to a model-based, computational perspective is change in statistical scientific thought at a more fundamental level. As researchers become increasingly involved in more complex stochastic model building enabled by advanced Bayesian computational methods, they also become more and more exposed to the inherent logic and directness of Bayesian model building. Scientifically relevant, highly structured stochastic models are often simply naturally developed from Bayesian formalisms and have overt Bayesian components. Hierarchical models with layers of random effects, random processes in temporal or spatial systems, and large-scale latent variables models of many flavours are just a few generic examples of nowadays standard stochastic structures in wide application, and that are all inherently Bayesian models. Much of the rapid growth in adoption of Bayesian methods from pragmatic viewpoints is engendering deeper, foundational change in scientific philosophy towards a more holistically Bayesian perspective. And this, in turn, has important implications for the core of the discipline; bringing Bayesian methods of stochastic modelling centre-stage – with models of increasing complexity and structure for reasons of increased realism – is inevitably re-energizing the core of the discipline, presenting new conceptual and theoretical challenges to Bayesian researchers as applied problems scale in dimension and complexity.

## The Handbook

*The Handbook of Applied Bayesian Analysis* is a showcase of contemporary Bayesian analysis in important and challenging applied problems, bringing together chapters contributed by leading researchers and practitioners in interdisciplinary Bayesian analysis. Each chapter presents authoritative discussions of currently topical application areas together with key aspects of Bayesian analysis in these areas, and takes the reader to the cutting edge of research in that topic. Importantly, each chapter is built around the application, and represents personal interests, experiences and views of the authors speaking from deep and detailed expertise and engagement in the applied problem area.

Each chapter of the *Handbook* involves a concise review of the application area, describes the problem contexts and goals, discusses aspects of the data and overall statistical issues, and develops detailed analysis with relevant Bayesian models and methods. Discussion generally contacts current frontiers of research in each application, with authors presenting their own perspectives, their own latest thinking, and highlighting their own research in both the application and in related and relevant Bayesian methodology used in their application. Each chapter also includes references linking to core publications in the applied field as well as relevant models and computational methods, and some also provide access to data, software and additional material related to the study of the chapter.

Importantly, each chapter contains appendix material that adds further foundational and supporting discussion of two flavours: material on the basic statistical models and methods, with background and key references, for readers interested in going further into methodological aspects, and more traditional appendix material representing additional technical developments in the specific application. Collectively, the appendices are an important component and distinctive feature of the *Handbook*, as they reflect a broad selection of models and computational tools used widely in applied Bayesian analysis across the diverse range of applied contexts represented.

## Chapter outlines

Chapters are grouped by broad field of application, namely

- Biomedical and Health Sciences
- Industry, Economics and Finance
- Environment and Ecology
- Policy, Political and Social Sciences
- Natural and Engineering Sciences

Inevitably selective in terms of broad fields as well as specific application contexts within each broad area, the chapters nevertheless represent topical, challenging and statistically illuminating studies in each case. Chapters within each area are as follows.

## Biomedical and Health Sciences

*Dunson* discusses an epidemiological study involving pregnancy outcomes. This chapter showcases Bayesian analysis in epidemiological studies that collect continuous health outcomes data, and in which the scientific and clinical interest typically focuses on the relationships between exposures and risks of an abnormal response, corresponding to an observation in the tails of the distribution. As there is minimal interest in relationships between exposures and the centre of the response distribution in such studies, traditional regression models are inadequate. For this reason, epidemiologists typically categorize both the outcome and the predictors, with the resulting inferences very sensitive to this categorization. Bayesian analysis using density regression, mixtures and nonparametric models, as developed and applied in this pregnancy outcome study, avoid and overcome these challenges.

*Green, Mardia, Nyirongo and Ruffieux* discuss the alignment of biomolecules. This chapter showcases Bayesian methods for shape analysis to assist with understanding the three-dimensional structure of protein molecules, which is one of the major unsolved biological challenges. This chapter addresses the problem of matching instances of the same structure in the CoMFA (Comparative Molecular Field Analysis) database of steroid molecules, where the three-dimensional coordinates of all the atoms in each molecule are stored. The matching problem is challenging because two instances of the same three-dimensional structure in such a database can have very different sets of coordinates, due not just to noisy measurements but also to rotation, translation and scaling. The authors present an efficient Bayesian methodology to identify, given two or more biomolecules represented by the coordinates of their atoms, subsets of those atoms which match within measurement error, after allowing for appropriate geometrical transformations to align the biomolecules.

*Cheng and Madigan* discuss a study of pharmaceutical testing from multiple clinical trials concerned with side-effects and adverse events among patients treated with a popular pain-relieving drug. This chapter showcases the development of sensitive Bayesian analysis of clinical trials studies involving problems of missing data and particularly non-ignorable dropout of patients from studies, as well as sequential methods and meta-analysis of multiple studies. The study concerns Vioxx, an anti-inflammatory drug that was licensed for use in the

USA by the FDA in 1999, and then withdrawn from the market in 2004 due to cardiovascular safety concerns. Merck, the manufacturer of Vioxx, conducted many clinical trials both before and after 1999. In part to avoid potential future scenarios like Vioxx, analyses of the data from these multiple clinical trials are of considerable importance and interest. The study raises multiple, challenging statistical issues and questions requiring sensitive evaluation, and the chapter highlights the utility of Bayesian analysis in addressing these challenges.

*Oakley and Clough* discuss uncertainty in a mechanistic model that has been used to conduct a risk assessment of contamination of farm-pasteurized milk with the bacterium Vero-cytotoxigenic *E. coli* (VTEC) O157. This chapter showcases Bayesian methods for analysing uncertainties in complex computer models. The VTEC model has uncertain input parameters, and so outputs from the model used to inform the risk assessment are also uncertain. The question then arises of how to reduce output uncertainty most efficiently. The authors conduct a variance-based sensitivity analysis to identify the most important uncertain model inputs, and so prioritize what further research would be needed to best reduce model output uncertainty.

*Schmidt, Hoeting, Pereira and Vieira* discuss temporal prediction and spatial interpolation for out-breaks of malaria over time for municipalities in the state of Amazonas, Brazil. This chapter showcases Bayesian spatial-temporal modelling for epidemiological discrete count data. Malaria is a world-wide public health problem with 40% of the population of the world at risk of acquiring the disease. It is estimated that there are over 500 million clinical cases of malaria each year world-wide. This work falls in the area of disease mapping, where data on aggregate incidence of some disease is available for various administrative areas, but the data for Amazonas are incomplete, covering only a subset of the municipalities. Furthermore, the temporal aspect is important because malaria incidence is not constant over time. A free-form spatial covariance structure is adopted which allows for the estimation of unobserved municipalities to draw on observations in neighbouring areas, but without making strong assumptions about the nature of spatial relations. A multivariate dynamic linear model controls the temporal effects and facilitates the forecasting of future malaria incidence.

*Merl, Lucas, Nevins, Shen and West* discuss a study in cancer genomics. This chapter showcases the application of Bayesian concepts and methods in an overall strategy for linking the results of *in vitro* laboratory studies of gene expression to *in vivo* human observation studies. The basic problem of translating inferences across contexts constitutes a generic, critical, and growing challenge in modern biology, which typically moves from laboratory experiments with cultured cells, to animal model experiments, to human outcome studies and clinical trials. The study described here concerns this problem in



the context of the genomics of several oncogene pathways that are fundamental to many human cancers. The application involves Bayesian sparse multivariate regression and sparse latent factor models for large-scale multivariate data, and details the use of such models to define and relate statistical signatures of biological phenomena between contexts. In addition, the study requires linking the resulting, model-based inferences to known biology; this is achieved using Bayesian methods for mapping summary inferences to databases of biological pathways. The study includes detailed discussion of biological interpretations of experimentally defined gene expression signatures and their elaborated subsignature representations emerging from Bayesian factor analysis of *in vivo* data, model-generated leads to design new biological experiments based on some of the findings, and contextual discussions of connections to clinical cancer profiling and prognosis.

*Henderson, Boys, Proctor and Wilkinson* discuss oscillations observed in the levels of two proteins, p53 and Mdm2, in single living cancer cells. This chapter showcases Bayesian methods in systems biology using genuine prior information and MCMC computation. The p53 tumour suppressor protein plays a major role in cancer. It has been described as the 'guardian of the genome', blocking cell cycle progression to allow the repair of damaged DNA. An increase of p53 due to stress causes an increase in the level of Mdm2 which in turn inhibits p53. From observations of levels of these two proteins in individual cancer cells, the objective is to learn about the rate parameters that control this feedback loop. By examining several cells, it is hoped to understand why they do not oscillate in phase. However, the modelling of the complex reactions within the cell makes this exercise highly computationally intensive. The authors develop a Bayesian approximation to the discrete time transition probabilities in the underlying continuous time stochastic model. Prior information about the unknown rate parameters is incorporated based on experimental values in the literature and they apply sophisticated MCMC methods to compute posterior distributions.

*Dawid, Mortera and Vicard* discuss the problem of evaluating the probability of a putative father being the real father of a child, based on his DNA profile and those of the mother and child. The chapter is a showcase for careful probabilistic reasoning. In recent years there has been heavy media coverage of DNA profiling for criminal identification, but the technique has also been useful in illuminating a number of complex genetic problems, including cases of disputed paternity. The paternity problem is complicated by the possibility of genetic mutation: the putative father could be the real father, yet a mutation in the child's DNA could seem to imply that he is not. The probability of paternity now depends strongly on the rate of mutation. On the other hand, estimates of mutation rates are themselves very sensitive to assumptions about paternity.

Using Austrian-German casework data, the authors present a meticulous study of this problem, constructing and analysing a model to handle paternity and mutation jointly.

## Industry, Economics and Finance

*Popova, Morton, Damien and Hanson* discuss a study in Bayesian analysis and decision making in the maintenance and reliability of nuclear power plants. The chapter showcases Bayesian parametric and semiparametric methodology applied to the failure times of components that belong to an auxiliary feedwater system. This system supplies cooling water during an emergency operation or to an electro-hydraulic control system, used for the control of the main electrical generating steam turbine. The parametric models produce estimates of the hazard functions that are compared to the output from a mixture of Polya trees model. The statistical output is used as the most critical input in a stochastic optimization model which finds the optimal replacement time for a system that randomly fails over a finite horizon. The chapter also discusses decision analysis, using the model in defining strategies that minimize expected total and discounted cost of nuclear plant maintenance.

*Cumming and Goldstein* discuss analysis of the Gullfaks oil field using a reservoir simulation model run at two different levels of complexity. This chapter showcases Bayes linear methods to address highly complex problems for which the full Bayesian analysis may be computationally intractable. A simulator of a hydrocarbon reservoir represents properties of the reservoir on a three-dimensional grid. The finer this grid is, the more accurately the simulator is expected to predict the real reservoir behaviour, but finer resolution also implies rapidly escalating computation times. Observed behaviour of the reservoir can in principle be used to learn about values of parameters in the simulator, but this Bayesian calibration demands that the simulator can be run many times at different values of these parameters in order to search for regions of parameter space in which acceptable matches are found to the observed data. The authors employ many runs of the simulator at low resolution to augment a few runs of the fine simulator. Their approach involves careful modelling of the relationship between the two versions of the simulator, as well as how the fine simulator relates to reality.

*Pievatolo and Ruggeri* discuss a study in Bayesian reliability analysis concerning underground train door failures in a European underground system over a period of nine years. The chapter showcases development and application of Bayesian stochastic process models in a reliability context. Facing questions about relevant reliability 'time' scales, the authors develop a novel bivariate Poisson process as a natural way to extend the usual Poisson models

for the occurrence of failures in repairable systems; the bivariate model uses both calendar time and kilometres driven by trains as metrics. An important consequence of this choice is that seasonal effects are easily incorporated into the model. The Bayesian models and MCMC methods developed lead to predictive distributions for failures and address key practical questions of how to assess reliability before warranty expiration, combining the data from several trains. This study also clarifies the advantages and disadvantages of using Poisson process models for repairable systems with a number of different failure modes.

*Ferreira, Bertolde and Holan* discuss an economic study of agricultural production in Espírito Santo State, Brazil, from 1990 to 2005. The chapter showcases the use of Bayesian multiscale spatio-temporal models that uses the natural geopolitical division of Espírito Santo State at levels of macroregions, microregions, and counties. The models involve multiscale latent parameters that evolve through time over a period of several years of the economic study. The analysis sheds light on the similarities and differences in agricultural production between regions within each scale of resolution, and on the temporal changes in relative agricultural importance of those regions as explicitly described by the evolution of the estimated multiscale components. The study involves a number of other questions relevant to the underlying spatio-temporal agricultural production process at each level of resolution, and builds on advanced Markov chain Monte Carlo methods for multivariate dynamic models integrated in an overall, highly-structured multiscale spatio-temporal framework.

*Lopes and Polson* discuss financial time series at the time of the 2007–08 credit crisis, showcasing the ability of Bayesian modelling and inference to represent a period of financial instability and to identify the underlying mechanisms. The authors consider several forms of model that exhibit stochastic volatility so as to capture the rapidly changing behaviour of the financial indicators at that time. Using Bayesian sequential model choice techniques, they show how the evidence accumulates over time that the pure stochastic volatility model is inferior to a model with jumps. Their work has implications for analysis and prediction in times of unusual market behaviour.

*Quintana, Carvalho, Scott and Costigliola* discuss studies in applications of the Bayesian approach to risk modelling regarding speculative trading strategies in financial futures markets. The chapter showcases applied Bayesian thinking in the context of financial investment management, highlighting the corresponding concepts of betting and investing, prices and expectations, and coherence and arbitrage-free pricing. Covering central applied methods and tools of Bayesian decision analysis and speculation in portfolio studies, risk modelling, dynamic linear models and Bayesian forecasting, and highly structured Bayesian graphical modelling approaches for multivariate, time-varying covari-

ance matrices in multivariate dynamic models, the chapter develops studies of investment strategies and returns in futures markets over a period between 1990 and 2008 based on portfolios of currency exchange rates, government bonds and stock market indices.

*Fernández-Villaverde, Guerrón-Quintana and Rubio-Ramírez* discuss macroeconomic studies of the dynamics of the US economy over the last 50 years using Bayesian analysis of dynamic stochastic equilibrium models. This chapter is a showcase of modern, model-based Bayesian analysis in mainstream economics studies, and an approach increasingly referred to as the *new macroeconomics*. The authors formulate and estimate a benchmark dynamic stochastic equilibrium model that captures much of the time-varying structure in the US macroeconomy over these years, and describe its application in policy analysis for public institutions – such as central banks – and private organizations and businesses. Application involves likelihood evaluations that are enabled using Bayesian sequential Monte Carlo and MCMC methods. The study discusses critical questions of the roles of priors and pre-sample information, documents a range of real and nominal rigidities in the US economy and discusses the increasingly central roles of such Bayesian approaches in this context as well as frontier research issues.

## Environment and Ecology

*Challenor, McNeall and Gattiker* discusses the potential collapse of the meridional overturning circulation in the Atlantic Ocean. This chapter showcases Bayesian methods for analysing uncertainty in complex models, and in particular for quantifying the risk of extreme outcomes. While climate science has concentrated on predictions of global warming, there are possible scenarios which, although with low probability, would have high impact. One such event is the collapse of the ocean circulation that currently ensures that Western Europe enjoys a warmer climate than, for instance, similar latitudes in Western North America. Collapse of the meridional overturning circulation (MOC) is predicted by the GENIE-1 climate model for some values of the model inputs, but the actual values of these inputs are unknown. A single run of GENIE-1 takes several hours, and the authors use Bayesian emulation to estimate the probability of MOC collapse based on a limited number of model runs, and to incorporate data comprising a sparse time series of five measurements of the MOC from 1957 to 2004.

*Clark, Bell, Dietze, Hersh, Ibanez, LaDeau, McMahon, Metcalf, Moran, Pangle and Wolosin* discuss demography of plant populations, showcasing applied Bayesian analysis and methods that allow for synthesis of information from multiple sources to estimate the demographic rates of trees and how they

respond to environmental variation. Data come from individual (tree) measurements over a period of 18 years, including diameter, crown area, maturation status, and survival, and from seed traps, which provide indirect information on fecundity. Different observations are available for different years and trees. The multiple data sets are synthesized with a process model where each individual is represented by a multivariate state-space submodel for both continuous (fecundity potential, growth rate, mortality risk, maturation probability) and discrete states (maturation status). Each year, state variables respond to a dynamic environment. Results provide unprecedented detail on the ways in which demographic rates relate to one another, within individuals over time, among individuals, and among species. The chapter also describes how results of these Bayesian methods are being used to assess how forests can respond to changing climate.

*Gelfand and Sahu* discuss environmental studies that aim to combine monitoring data and computer model outputs in assessing environmental exposure. This chapter showcases Bayesian data fusion methods using spatial Gaussian process models in studies of weekly deposition data from multiple US sites monitored by the US National Atmospheric Deposition Program. Environmental exposure community numerical models are now widely available for a number of air pollutants. Based on inputs from a number of factors such as meteorological conditions, land usage, and power station emission volumes, all of which are responsible for producing air pollution, and some predictions of spatial surfaces for current, past, and future time periods, these models provide output exposures at various spatial and temporal resolutions. For large spatial regions such as the entire United States, the spatial coverage of the available network monitoring stations can never match the coverage at which the computer models produce their output. However, the monitoring data will be more accurate than the computer model output since, up to measurement error, they provide the actual true levels: observations from the realization of the pollution process surface at that time. It is important to combine these two sets of information to make inference regarding pollution exposure, and this study represents best-Bayesian practices in addressing this problem.

*Choy, Murray, James and Mengersen* discuss eliciting knowledge from ecological experts about the habitat of the Australian brush-tailed rock-wallaby. This chapter is a showcase of techniques for eliciting expert judgement about complex uncertainties. The rock-wallaby is an endangered species, and in order to map where it is likely to be found, it is essential to use expert judgement about how the various environmental factors (such as geology, land cover and elevation) influence the probability of rock-wallabies being present at a site. The authors employ an indirect elicitation method in which the experts are presented with descriptions of some specific sites and asked for their probabilities. The relationship between probability of occurrence and the environmental

variables is then inferred and used to predict the rock-wallaby's likely habitats throughout the region of interest.

*Tebaldi and Smith* discuss studies in characterizing the uncertainty of climate change projections, showcasing Bayesian methods for integration and comparison of predictions from multiple models and groups. The chapter describes a suite of customised Bayesian hierarchical models that synthesize ensembles of climate model simulations, with the aim of reconciling different future projections of climate change, while characterizing their uncertainty in a rigorous fashion. Posterior distributions of future temperature and/or precipitation changes at regional scales are obtained, accounting for many peculiar data characteristics, such as systematic biases, model-specific precisions, region-specific effects, changes in trend with increasing rates of greenhouse gas emissions, and others. The chapter expands on many important issues characterizing model experiments and their collection into multimodel ensembles, and addresses the need of 'impact research', by proposing posterior predictive distributions as a representation of probabilistic projections. In addition, the calculation of the posterior predictive distribution for a new set of model data allows a rigorous cross-validation approach to assess and, in this study, confirm the reasonableness of the Bayesian modelling assumptions.

## Policy, Political and Social Sciences

*Carvalho and Rickershauser* discuss a study of temporal volatility and information flows in political campaigns, showcasing Bayesian analysis in evaluation of information impact on vote sentiment and behaviour in highly publicized campaigns. The core application is to the 2004 US presidential campaign. The study builds a measure of information flow based on the returns and volume of the 'Bush wins the popular vote in 2004' futures contract on the tradesports/intrade prediction market. This measure links events to information level, providing a direct way to evaluate its impact in the election. Among the findings are that information flows increased as a result of the televised debates, Kerry's acceptance speech at the Democratic convention, and national security-related stories such as the report that explosives vanished in Iraq under the US's watch, the CBS story about Bush's National Guard service and the subsequent retraction, and the release of the bin Laden tape a few days before the election. Contrary to popular accounts of the election, ads attacking Kerry's military service aired by the Swift Boat Veterans for Truth in August apparently contributed only a limited amount of information to the campaign. This political science application develops novel hidden state-space models of volatility in information flows and model fitting and evaluation using Bayesian MCMC methods for nonlinear state-space models.



*Gamerman, Soares and Gonçalves* discuss whether cultural differences may affect the performance of students from different countries in the various test items which make up the international PISA test of mathematics ability. This chapter showcases a Bayesian model that incorporates this kind of differential item functioning (DIF) and the role of prior information. The PISA tests in mathematics and other subjects are widely used to compare the educational attainment of 15-year old students in different countries; in 2009, 67 countries have taken part from around the world. DIF is a significant issue with the potential to compromise such comparisons between countries, and substantial DIF may remain in the administered test despite preliminary screening of candidate test items. The authors seek to discover the extent of DIF remaining in the mathematics test of 2003. They employ a hierarchical three-parameter logistic model for the probability of a correct response on an individual item, where the three parameters control the difficulty of the item, its discriminating power and its guessability, and their model allows for different kinds of DIF where any of these parameters may vary between countries. The authors' Bayesian model avoids identifiability problems faced by competing approaches and requires weaker hypotheses due, especially, to the important role played by the prior distributions.

*Heiner, Kennedy and O'Hagan* discuss auditing of the operation of the food stamps welfare scheme in the state of New York, USA, highlighting the power of Bayesian methods in analysing data that evolve over time. Auditors examine a sample of individual awards of food stamps to see if the value awarded is correct according to the rules of the scheme. The food stamps program is a federal scheme, and if a state is found to have too large an error rate in administering it the federal government can impose large financial penalties. In New York state, the program is administered by individual counties, and sizes of audit samples in small counties can be so small that only one or two errors are found in any given year. The authors propose a model that includes a nonparametric component for the error magnitudes (taints), a hierarchical model for overall error rates across counties and parameters controlling the variation of rates from one year to the next, including an overall trend in error rates. The model allows in particular for estimation of rates in small counties to be smoothed across counties and through time.

*Rubin, Wang, Yin and Zell* discuss a study in estimating the effects of 'treating hospital type' on cancer survival, using administrative data from central and northern Sweden via the Karolinska Institute in Stockholm. The chapter represents a showcase in application of Bayesian causal inference, in particular using the posterior predictive approach of the 'Rubin causal model' and methods of principal stratification. The central applied question, inferring which type of hospital (e.g. large patient volume versus small volume) is superior for treating certain serious conditions, is a difficult and important problem in institutional

assessment and comparisons in a medical context. Ideal data from randomized experiments are simply not available, leading to reliance on observational data. The study involves questions of which factors may reasonably be considered ignorable in the context of covariates available, and non-compliance complications due to transfers between hospital types for treatment, and showcases Bayesian causal modelling utilizing simulation-based imputation techniques.

## Natural and Engineering Sciences

*Cemgil, Godsill, Peeling and Whiteley* discuss musical audio signal analysis in the context of an application to multipitch audio and determining a musical ‘score’ representation that includes pitch and time duration summary for a musical extract (the so-called ‘piano-roll’ representation of music). This chapter showcases applied Bayesian analysis in audio signal processing in real environments where acoustical conditions and sound sources are highly variable, yet audio signals possess strong statistical structure. There is typically much prior information about underlying structures and the detail of the recorded acoustical waveform (physical mechanisms by which sounds are generated, cognitive processes by which sounds are perceived by the human auditory system, mechanisms by which high-level sound structures are compiled). A range of Bayesian hierarchical models – involving both time and frequency domain dynamic models, and methods of fitting using simulation-based and variational approximations – are developed in this chapter. The resulting models possess complex statistical structure and so highly adaptive and powerful computational techniques are needed to perform inference, as this study exemplifies.

*Higdon, Heitmann, Nakhleh and Habib* discuss perhaps the grandest of all problems, the nature and evolution of the universe. This chapter showcases techniques for emulating complex computer models with many inputs and outputs. The  $\Lambda$ -cold dark matter model is the simplest cosmological model in agreement with the cosmic microwave background and large scale structure measurements. This model is determined by a small number of parameters which control the composition, expansion and fluctuations of the universe, and the objective of this study is to learn about the values of these parameters using measurements from the Sloan Digital Sky Survey (SDSS). Model outputs include a dark matter spectrum for the universe and a temperature spectrum for the cosmic microwave background. A key component of the Bayesian analysis is to find a parsimonious representation of such high-dimensional output. Another is innovative modelling to combine the evidence from data on both



spectra to find which model input parameters influence the output appreciably, and to learn about those parameters.

*Liang, Jordan and Klein* discuss the use of probabilistic context-free grammars in natural language processing, involving a large-scale natural language parsing task. The chapter is a showcase of detailed, highly-structured Bayesian modelling in which model dimension and complexity responds naturally to observed data, building on the adaptive nature of the underlying nonparametric Bayesian models developed by the authors. The framework involves structured hierarchical Dirichlet process modelling and customized model fitting via variational methods, to address the core problem of identifying appropriate levels of model complexity in using probabilistic context-free grammars as important components in the modelling of syntax in natural language processing. Detailed development and evaluation in experiments with a synthetic grammar induction task complement the application to a large-scale natural language parsing study on data from the *Wall Street Journal* portion of the Penn Treebank, a large data set used in the natural language processing community for evaluating parsers.

*Lee, Taddy, Gramacy and Gray* discuss the development of circuit devices, bipolar junction transistors, which are used to amplify electrical current, and showcases the use of a flexible kind of emulator based on a treed Gaussian process. To aid with the design of the circuit device, a computer model predicts its peak output as a function of the input dosage and a number of design parameters. The peak output response can jump sharply with only small changes in dosage, a feature that the treed Gaussian process emulator is able to capture. The methodology also involves a novel sequential design procedure to generate data to fit the emulator, and performs sensitivity analysis and both calibration and validation using experimental data.

*Prado* discusses a study of experimental data involving large-scale EEG time series generated on individuals subject to tasks inducing cognitive fatigue, with the eventual goals of models able to predict cognitive fatigue based on non-invasive scalp monitoring of real-time EEG fluctuations. The chapter showcases the development and application of structured, multivariate Bayesian dynamic models for analysis of time-varying, non-stationary and erratic (brain wave) time series. Novel time-varying autoregressive and regime switching models, incorporating substantively relevant prior information via structured priors and fitted using novel, customized Bayesian computational methods, are described. The applied study involves an experimental subject asked to perform simple arithmetic operations for a period of three hours. Prior to the experiment, the subject was confirmed to be alert. After the experiment ended, the subject was fatigued, as determined by measures of performance and post-task mood. The study shows how the Bayesian analysis is used to assist practitioners in real time detection of cognitive fatigue.