Preface

When we perceive the physical world, make a decision, and take an action, a critical issue that our brains must deal with is uncertainty: there is uncertainty associated with the sensory system, the motor apparatus, one's own knowledge, and the world itself. The Bayesian framework of statistical estimation provides a coherent way of dealing with these uncertainties. Bayesian methods are becoming increasingly popular not only in building artificial systems that can handle uncertainty but also in efforts to develop a theory of how the brain works in the face of uncertainty.

At the core of the Bayesian way of thinking is the Bayes theorem, which maintains, in its simplest interpretation, that one's belief about the world should be updated according to the product of what one believed in before and what evidence has come to light since. The strength of the Bayesian approach comes from the fact that it offers a mathematically rigorous computational mechanism for combining prior knowledge with incoming evidence.

A classical example of Bayesian inference is the Kalman filter, which has been extensively used in engineering, communication, and control over the past few decades. The Kalman filter utilizes knowledge about the noise in sensory observations and the dynamics of the observed system to keep track of the best estimate of the system's current state and its variance. Although Kalman filters assume linear dynamics and Gaussian noise, recent Bayesian filters such as particle filters have extended the basic idea to nonlinear, non-Gaussian systems. However, despite much progress in signal processing and pattern recognition, no artificial system can yet match the brain's capabilities in tasks such as speech and natural scene recognition. Understanding how the brain solves such tasks could offer considerable insights into engineering artificial systems for similar tasks.

A Bayesian approach can contribute to an understanding of the brain at multiple levels. First, it can make normative predictions about how an ideal perceptual system combines prior knowledge with sensory observations, enabling principled interpretations of data from behavioral and psychophysical experiments. Second, algorithms for Bayesian estimation can provide mechanistic interpretations of neural circuits in the brain. Third, Bayesian methods can be used to optimally decode neural data such as spike trains. Lastly, a better understanding the brain's computational mechanisms should have a synergistic impact on the development of new algorithms for Bayesian computation, leading to new applications and technologies.

About This Book

This book is based on lectures given at the First Okinawa Computational Neuroscience Course, held in November 2004 at Bankoku Shinryokan, Okinawa, Japan. The intention of the course was to bring together both experimental and theoretical neuroscientists employing the principles of Bayesian estimation to understand the brain mechanisms of perception, decision, and control.

The organization of the book is as follows. In the Introduction, Doya and Ishii give the mathematical preliminaries, including the Bayes theorem, that are essential for understanding the remaining chapters of the book. The second part of the book, Reading Neural Codes, introduces readers to Bayesian concepts that can be used for interpretation of neurobiological data. The chapters by Fairhall, Pillow, and Richmond and Wiener explore methods for characterizing what a neuron encodes based on its spike trains. The chapter by Penny and Friston describes how Bayesian theory can be used for processing and modeling functional brain imaging data. The third part, entitled Making Sense of the World, assembles chapters on models of sensory processing. Pouget and Zemel review ideas about how information about the external world can be coded within populations of neurons. Latham and Pouget explore the use of such codes for neural computation. Lee and Yuille consider how top-down and bottom-up information can be combined in visual processing, while Knill uses Bayesian models to investigate optimal integration of multiple sensory cues. The final part, Making Decisions and Movements, explores models of the dynamic processes governing actions and behaviors. The chapter by Shadlen, Hanks, Churchland, Kiani, and Yang focuses on neurons in higher visual cortex that accumulate evidence for perceptual decisions over time. Rao discusses a model of how cortical circuits can implement "belief propagation," a general method for Bayesian estimation. Todorov reviews optimal control theory from the viewpoint of Bayesian estimation. Finally, Körding and Wolpert utilize Bayesian decision theory to understand how humans make decisions about movements.

It is our hope that this book will stimulate further research into Bayesian models of brain function, leading to a deeper, mathematically rigorous understanding of the neural processes underlying perception, decision, and action.

Acknowledgments

We wish to thank The Cabinet Office, Japanese Neural Network Society, Tamagawa University, and Kyushu Institute of Technology for sponsoring Okinawa Computational Neuroscience Course 2004, and Sydney Brenner, President of Okinawa Institute of Science and Technology, for advocating the course.

The course and the subsequent book publication would have been impossible without all the secretarial work by Izumi Nagano. Emiko Asato and Yasuhiro Inamine also helped us in manuscript preparation. Our thanks go to MIT Press editors Barbara Murphy and Kate Blakinger and to Terry Sejnowski for supporting our book proposal.

AP would like to thank the National Science Foundation (NSF) for their support. RPNR would like to acknowledge the support of the National Science Foundation (NSF), the ONR Adaptive Neural Systems program, the Sloan Foundation, and the Packard Foundation.