Preface

Nonlinear time series methods have developed rapidly over a quarter of a century and have reached an advanced state of maturity during the last decade. Implementations of these methods for experimental data are now widely accepted and fairly routine, however genuinely useful applications remain rare. The aim of this book is to focus on the practice of applying these methods to solve real problems. It is my hope that the methods presented here are sufficiently accessible, and the examples sufficiently detailed, that practitioners in other areas may use this work to begin considering further applications of nonlinear time series analysis in their own disciplines.

This volume is therefore intended to be accessible to a fairly broad audience: both specialists in nonlinear time series analysis (for whom many of these techniques may be new); and, scientists in other fields (who may be looking to apply these methods within their speciality). For the experimental scientist looking to use these methods, MATLAB implementation of the underlying algorithms accompany this book.

Although the mathematical motivation for nonlinear time series analysis is fairly advanced, I have chosen to keep technical content in this book to a minimum. Postgraduate and advanced undergraduate students in the physical sciences should find the material reasonably easy to understand. This book may be read sequentially; skimmed in a pseudo-random order; used primarily as a reference; or, treated as a manual for the companion computer programs.

The applications

To illustrate the usefulness of nonlinear time series analysis, a wide variety of physical, financial and physiological systems have been considered.

In particular, several detailed applications serve as case studies of fruitful (and occasionally less fruitful) applications, and illustrate the mathematical techniques described in the text. These applications include:

- diagnosis and control of cardiac arrhythmia in humans (prediction of ventricular fibrillation);
- characterisation and identification of aberrant respiratory control in infants (sleep apnea and Sudden Infant Death syndrome);
- simulation and recognition of human vocalisation patterns;
- interpretation and prediction of financial time series data; and,
- quantitative assessment of the effectiveness of government control measures of the recent SARS crisis in Hong Kong.

The applications described in this book are drawn largely from my own work, and for this I offer no apology. These applications are the applications which interest me most, and with which I am most familiar. Some of the applications are particularly long and detailed: this is so that the reader can truly get a feel for the complexity of each example and the effort required to fit the methods to the problems. However, if one is primarily interested in the methods themselves, it is entirely possible to read this monograph as a textbook: one should simply skip (or skim) the long-winded application based sections.

The tools

The technical tools utilised in this book fall into three distinct, but interconnected areas: quantitative measures of nonlinear dynamics, Monte-Carlo statistical hypothesis testing (aka Surrogate data analysis), and nonlinear modelling.

In Chapter 1 we discuss nonlinear time series analysis from the perspective of time delay embedding reconstruction. We describe the current state-of-the-art in reconstruction techniques, together with my own view of how to "best" embed scalar time series data. Chapters 2 and 3 are concerned with the estimation of statistical quantities (mostly dynamic invariants) from time series data.

Quantitative measures, such as estimation of correlation dimension and Lyapunov exponents have been described previously in standard texts. However, this material is included here for completeness: an understanding of these techniques is necessary for the surrogate data methods that follow.

More importantly, it is necessary to provide a current description of the most appropriate estimation techniques for short and noisy data sets. In particular, it should be noted that many of the more venerable techniques perform poorly when faced with the limitations of most experimentally obtained data i.e. short and noisy time series. Hopefully, some of the new methods we discuss will help alleviate these problems.

Three standard methods of Monte-Carlo statistical hypothesis testing (the method of surrogate data) have come into widespread use recently. In Chapters 4 and 5, these methods are described along with their many limitations and restrictions. Several new, state-of-the-art methods to circumvent these problems are also described and demonstrated to have useful application. Finally, a natural extension of surrogate data methods is the testing of model fit, and this is the third focus of this book.

Standard nonlinear modelling methods (neural networks, radial basis functions and so on) are the subject of numerous excellent texts. In Chapter 6 we focus on finding the best model and how to determine when a given model is "good enough". These are two problems that are not well addressed in the current literature. For this work, we utilise the surrogate data methods described previously, as well as information theoretic measures of both the model and the data. These information theoretic measures of the data are in turn related to, and form the foundation of, the dynamic invariants described in the introduction of this book.

Acknowledgments

This volume is actually the result of many years work. The new methods presented here build on a broad and strong foundation of nonlinear time series analysis and nonlinear dynamical systems theory. In collating this material in a single volume I should thank many people. I would never have met this exciting research field if it were not for my PhD adviser, Dr. Kevin Judd (University of Western Australia, Perth), and also Dr. Steven Stick (Princess Margaret Hospital for Children, Perth). My interest in the analysis of human ECG rhythms was sown by Prof. Robert Harrison (Heriot-Watt University, Edinburgh), and my enthusiasm for financial time series analysis is due to the significant work of Prof. Marius Gerber (Proteus VTS). I should wish also acknowledge Prof. Michael Tse (Hong Kong Polytechnic University) for his help and also for suggesting this endeavour in the first place: without his incessant encouragement I would not have

written this.

As is usual for a post-doctoral researcher, my work has been conducted in a variety of settings. Primarily, the work that serves as the foundation for this volume was conducted in the Centre for Applied Dynamics and Optimisation (CADO) in the Mathematics Department of the University of Western Australia; the Nonlinear Dynamics Group in the Physics Department of Heriot-Watt University (Edinburgh); and, the Applied Nonlinear Circuits and Systems Research Group in the Department of Electronic and Information Engineering at the Hong Kong Polytechnic University. I extend my warm thanks to my numerous friends and colleagues in all three centres.

Finally, I would like to acknowledge those who have contributed directly to this work: Dr. Tomomichi Nakamura persevered through an early draft version of much of this manuscript and gave me valuable feedback at each stage, and my students Mr. Xiaodong Luo and Mr. Yi Zhao assisted and provided me with valuable ideas and new ways of looking at most of the content of this work. Most important of all, my wife Sylvia endured the entire text of this volume and vastly improved both the grammar and style.

Finally, I am obliged to acknowledge various sources of financial support. This work was supported in part by Hong Kong University Grants Council (UGC) Competitive Earmarked Research Grants (CERG) numbers PolyU 5235/03E and PolyU 5216/04E as well as a direct allocations from the Department of Electronic and Information Engineering (A-PE46 and A-PF95).

Feedback

Although the best way to really understand how these techniques work is to create your own implementation of the necessary computer code, this is not always possible. Therefore, I have made MATLAB implementations of many of the key algorithms and the various case studies are available from my website (`http://small.eie.polyu.edu.hk/`). You are most welcome to send me your comments and feedback (`ensmall@polyu.edu.hk`).

Michael Small
Hong Kong
January 25, 2005