Identifying Evolving Multivariate Dynamics in Individual and Cohort Time Series, with Application to Physiological Control Systems

by

Shamim Nemati

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2013

ARCHIVES
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
APR 0 2 2013
LIBRARIES

© Massachusetts Institute of Technology 2013. All rights reserved.

Author .....
Department of Electrical Engineering and Computer Science
December 27, 2012

Certified by .....
George C. Verghese
Professor
Thesis Supervisor

Certified by .....
Atul Malhotra
Associate Professor
Thesis Supervisor

Accepted by .....
Leslie A. Kolodziejski
Chairman, Department Committee on Graduate Theses

Identifying Evolving Multivariate Dynamics in Individual and Cohort Time Series, with Application to Physiological Control Systems

by

Shamim Nemati

Submitted to the Department of Electrical Engineering and Computer Science
on December 28, 2012, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

Physiological control systems involve multiple interacting variables operating in feedback loops that enhance an organism's ability to self-regulate and respond to internal and external disturbances. The resulting multivariate time-series often exhibit rich dynamical patterns, which are altered under pathological conditions. However, model identification for physiological systems is complicated by measurement artifacts and changes between operating regimes. The overall aim of this thesis is to develop and validate computational tools for identification and analysis of structured multivariate models of physiological dynamics in individual and cohort time-series.

We first address the identification and stability of the respiratory chemoreflex system, which is key to the pathogenesis of sleep-induced periodic breathing and apnea. Using data from both an animal model of periodic breathing, as well as human recordings from clinical sleep studies, we demonstrate that model-based analysis of the interactions involved in spontaneous breathing can characterize the dynamics of the respiratory control system, and provide a useful tool for quantifying the contribution of various dynamic factors to ventilatory instability. The techniques have suggested novel approaches to titration of combination therapies, and clinical evaluations are now underway.

We then study shared multivariate dynamics in physiological cohort time-series, assuming that the time-series are generated by switching among a finite collection of physiologically constrained dynamical models. Patients whose time-series exhibit similar dynamics may be grouped for monitoring and outcome prediction. We develop a novel parallelizable machine-learning algorithm for outcome-discriminative identification of the switching dynamics, using a probabilistic dynamic Bayesian network to initialize a deterministic neural network classifier. In validation studies involving simulated data and human laboratory recordings, the new technique significantly outperforms the standard expectation-maximization approach for identification of switching dynamics. In a clinical application, we show the prognostic value of assessing evolving dynamics in blood pressure time-series to predict mortality in a cohort of intensive care unit patients.

A better understanding of the dynamics of physiological systems in both health and disease may enable clinicians to direct therapeutic interventions targeted to specific underlying mechanisms. The techniques developed in this thesis are general, and can be extended to other domains involving multi-dimensional cohort time-series.

Thesis Supervisor: George C. Verghese
Title: Professor, EECS, MIT

Thesis Supervisor: Atul Malhotra
Title: Associate Professor, Medicine, Harvard

Acknowledgments

To my Parents,
For their Uncompromising Dedication to
the Cause of Education in the Face of Adversity.
And to Andre,
For his Pure Heart and Curious Mind!

I wish to acknowledge those individuals who have been an indispensable source of support to me throughout the years. First and foremost, I would like to acknowledge members of my committee for trusting me with my endless explorations, and providing the crucial guidance when needed: Professor **Atul Malhotra**'s friendly guidance, generous support, and constant encouragement have made this thesis possible, and I am forever grateful to him.
Professor **George Verghese** has been an excellent mentor and a truly compassionate individual. His legendary dedication to education and caring for students, at a very personal level, is one of the greatest treasures the EECS department at MIT has to offer.
Professor **Roger Mark** provided me with a research group and an office ("a second home"), when I most needed one. Roger's door was always open for help, whether it was for academic advising or for student's personal issues. Roger and Dottie were the most caring housemasters at the Sidney Pacific graduate residence, where I spent my first two years at MIT. Dr. Mark's passion for physiology inspired many students, including myself, to pursue a career in biomedical research.
Professor **James Butler** was simply inspirational! I learned more than a few things about "critical thinking" from him, which I will cherish for the rest of my life. Jim's expertise in pulmonary physiology has been invaluable to the design and implementation of several of the studies presented in this thesis.

I am very grateful to two very special individuals- **Farah Flaugher** and **Terry Flaugher**- my dear aunt and uncle, for their unconditional support and love; that I will treasure for the rest of my life! I also would like to thank my brother, **Shaya**, for being my greatest source of support as we embarked on a risky journey to the United States over a decade ago.

I would like to thank the many teachers and professors who have graciously provided me with their guidance and wisdom throughout my life:
I'm very grateful to my first teacher, Mr. **Askhari**, who taught me the Persian alphabet, and basic arithmetic. Upon our arrival in the United States, my brother and I both benefited greatly from the teachers and professors at the University of Central Oklahoma (UCO). In particular, Professor **Dan Endres** was a great source of encouragement and support. His mathematics competitions and encouragement set me on a path to majoring in mathematics as an undergrad. Professor **Mark Yeary**, at the University of Oklahoma (OU), was my first true academic role model. I was extremely privileged to have him as my mentor. With his encouragement and support, I wrote my first academic paper, and that set me on a path to pursue a career in academics. I am also very grateful to Professor **Murad Özaydin**, from the Department of Mathematics at OU, for mentoring me through a special postgraduate program in signal processing and computational and applied mathematics (SIGCAM). My most heartfelt gratitude goes to Dr. **Gari Clifford** for introducing me to the field of physiological signal processing. Gari was a highly enthusiastic teacher and a great insightful communicator. His willingness to sit down and discuss career paths and projects with students from all walks of life, and his philosophical disposition towards openness and sharing in scientific collaborations, made him an invaluable asset for MIT and a great mentor for the students in our lab and

beyond. I also would like to thank Professor **Ary Goldberger** and Dr. **Madalena Costa** (of Beth Israel Deaconess Medical Center) for their openness to discuss ideas and provide guidance. Again I'm grateful to Gari for introducing me to such amazing individuals in the MIT and Harvard arenas, and at scientific conferences. Professor **Ryan Adams**, Harvard School of Engineering, has been instrumental in shaping my research direction towards the last year of my PhD. I benefited greatly from his willingness to meet on a regular basis to discuss the application of machine-learning to computational physiology.

I also would like to acknowledge those colleagues without whose help this work could not have been possible:
I am grateful to Dr. **Thomas Heldt**, for kindly providing the tilt-table data and his willingness to explain difficult physiological concepts. I would also like to thank Dr. **Faisal Kashif** for his insightful discussions and his encouragements. Special thanks are due to Dr. **Brad Edwards** for patiently teaching me about respiratory physiology, and his openness to share data (in particular, for his amazing experiments on a lamb model of apnea). It has been a great privilege to work with Dr. Edwards, and to get to know him at a personal level. I learned so much from Brad and Dr. **Scott Sands** about how to prepare scientific manuscripts and communicate abstruse scientific ideas. I'm also very thankful to the other members of the Malhotra Lab at the Harvard Medical School, Drs. **Andrew Wellman, Robert Owens, Lisa Campana, Julian Saboisky**, and others. I would like to specially acknowledge **Pam DeYoung** for her encouraging words and acts of kindness. I have also truly enjoyed and benefited from the mathematics and neuroscience-related discussions with Drs. **Ben Polletta** and **Bob Schafer**, and look forward to continuing collaborations in the coming years.

I am very thankful to the members of the LCP group at MIT-HST: Drs. **Leo Celli**, **Joon Lee**, **Li-wei Lehman**, **Dan Scott**, **Ikaro Silva**, and **George Moody**, as well as, **Mauricio Villarroel** and **Ken Pierce**, for all the efforts they have put into the MIMIC and the Physionet databases, and for their philosophy of sharing and their readiness to provide help whenever needed. I have enormously benefited from all of their knowledge and expertise in Linux, Matlab, as well as their knowledge of sailing and life!

My friends have kept me going all these years, making my life more interesting and fun, and for that I am so grateful to:
Special thanks go to my dear friend, **Ibon Santiago**: "*Les amis sont des compagnons de voyage, qui nous aident a avancer sur le chemin d'une vie plus heureuse.*" Ibon's commitment to human beings, loving, and forgiving is one of the most precious sources of inspiration that I will take away from my time at MIT. **Igor Iwanek** was and always will be a breath of fresh air; thank you for all the Salsa outings! I would also like to take this opportunity to sincerely thank my other friends: **Mona Khabazan**, **Gina Marciano**, Dr. **Manjola Ujkaj**, Dr. **Cosette Chichirau**, Dr. **Hila Hashemi**, **Rene Lucena**, Dr. **Dimitris Baltzis**, **Kevin Brokish**, and so many other wonderful friends at MIT and Harvard whose names may have escaped me temporarily.

This research was partially made possible by the National Institutes of Health (NIH) T32 training grant (HL07901). The contents of this thesis are solely my responsibility and do not represent the views of NIH or any other legal entity.

"*Alas*", said the mouse, "*the whole world is growing smaller every day. At the beginning it was so big that I was afraid, I kept running and running, and I was glad when I saw walls far away to the right and left, but these long walls have narrowed so quickly that I am in the last chamber already, and there in the corner stands the trap that I must run into.*" "*You only need to change your direction,*" said the cat, and ate it up.
–Kafka little fable

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Motivations for Study of Dynamics in Physiological Systems

Feedback and adaptation are key characteristics of physiological systems, which are confronted by the task of operating under the influence of continuously changing internal and external factors. For instance, the human respiratory system has to adapt to extreme changes in metabolic demand that spans a wide spectrum between quiet sleeping and active exercise. To maintain adaptability, physiological systems have to integrate sensory inputs of varying modality and across multiple time-scales, to produce timely and appropriate responses. For example, afferent feedback to the respiratory centers in the brainstem include signals from the chest and lung stretch receptors (with delays of the order of a few hundred milliseconds), and the carotid body and the brainstem chemoreceptors (with delays of the order of a few seconds to minutes, respectively). As a consequence of complex feedback interactions among physiological variables, physiological systems exhibit rich dynamical patterns, ranging from periodic oscillations (for an example see Fig. 1-1), to rhythms with differing periodicities, to chaotic and noise-like fluctuations. These patterns are altered under pathological conditions, with the appearance of, for example, periodic oscillations at time-scales that are not typically present in normal subjects (e.g., Cheyne-Stokes breathing in congestive heart failure patients), or the loss of rhythmicity in normally rhythmic processes (e.g., atrial or ventricular fibrillation replacing normal sinus rhythm) [1].

A central premise of mathematical modeling in "systems physiology" is the idea that pathological dynamics result from an otherwise intact physiological control system operating within an undesirable region of its parameter space [1]. For instance, the periodic breathing pattern shown in Fig. 1-1 is generally attributed to a chemoreflex feedback loop with hypersensitive chemosen-

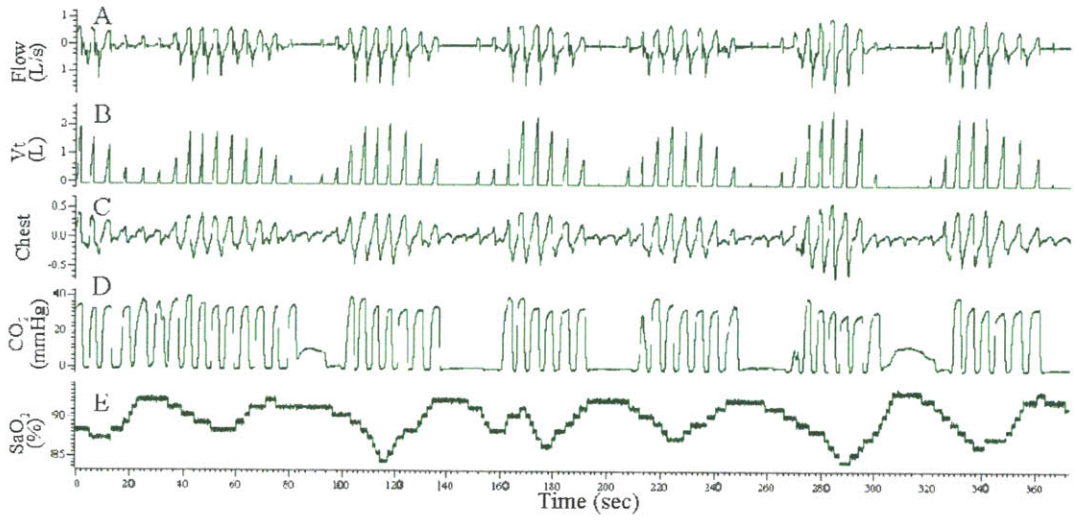

Figure 1-1: Sleep-induced periodic breathing in a human subject.
Measurements of respiratory flow (A), tidal volume (B), chest movements (C), partial pressure of CO_2 in exhaled air (D), oxygen saturation (E). The waxing and waning pattern of breathing followed by apnea, apparent in the tidal-volume signal, is accompanied by a drop in blood oxygen, rise in blood CO_2 concentration, and often an arousal at the end of each bout of apnea. These arousals are generally marked by an increase in sympathetic activity and an increase in blood pressure that often carries over to the daytime. The remarkable periodicity of the breathing patterns depicted above is generally attributed to a hypersensitive or unstable chemoreflex control system.

sors [2]. Therefore, understanding the mechanisms involved in healthy and unhealthy dynamics enables clinicians to direct therapeutic interventions at the specific underlying causes. A primary goal of systems physiology is to understand the physiological changes in system dynamics that occur as a result of administration of drugs, mechanical interventions in the function of organs, or due to physiological changes that accompany aging and disease. Therefore, given measurements of physiological variables, a system physiologist aims to identify the dynamics governing their interaction, for example to quantify the effects of interventions on unobserved system variables (e.g., "drug X enhances chemosensitivity"), or to predict system response to interventions (e.g., "patient Y will likely exhibit unstable breathing if exposed to a hypoxic gas mixture").

We now briefly discuss three illustrative examples of dynamical patterns in physiological time-series, which we will use as our motivation for development and validation of computational tools throughout this thesis.

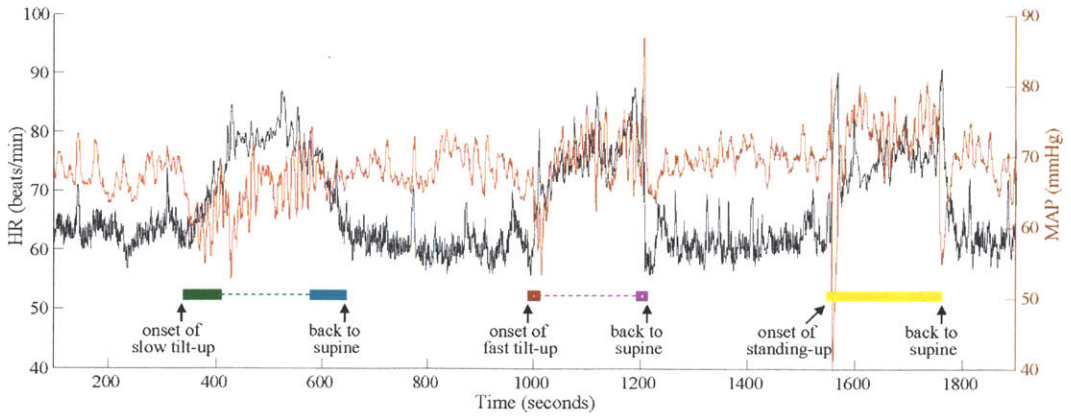

Figure 1-2: Dynamic regulation of heart rate and blood pressure.
Time series of of heart rate (HR; black) and mean arterial blood pressure (MAP; blue), measured at the level of the heart, are shown for a 30 minute recording from a tilt-table experiment (Heldt et al. [16]). The timings of three tilting events– slow tilt up and back to supine (green to cyan), fast tilt up and back to supine (red to magenta), and standing and back to supine (yellow)– are marked in sequence. Tilting is accompanied by a temporary drop in blood pressure, which gradually returns to its baseline value, due to the compensatory response of the baroreflex control system.

1.1.1 Chemoreflex Feedback Loop and Ventilatory Instability

As a first example, let us consider breathing instabilities in the form of *cyclic apnea* or *periodic breathing* (see Fig. 1-1). Periodic breathing is commonly observed in both preterm and term infants [3–6], as well as in adult subjects at altitude [7–10], and in patients with congestive heart failure [11–14]. Although the specific mechanisms underlying each condition may vary, ventilatory instability can result from increases in the ventilatory sensitivity to hypoxia/hypercapnia (*controller* gain), in the efficiency of gas exchange (*plant* gain), or in the circulatory delay between lungs and chemoreceptors [8, 12, 13, 15]; these together define the dynamic *loop gain* of the respiratory control system. The concept of dynamic loop gain has been increasingly emphasized because it integrates each of these physiological components into a single function of frequency that describes the stability of the respiratory control system in the presence of delays in the loop. High loop gain generally describes a system that is intrinsically less stable, whereas a low loop gain describes a more stable system. Analysis of the underlying mechanism involved in periodic breathing is crucial for elucidating the primary causes of instability, as well as predicting the response to therapies in individual patients.

1.1.2 Baroreflex Feedback Loop and Control of Heart Rate and Blood Pressure

In mammals, beat-to-beat values of heart rate (HR) and blood pressure (BP) are highly coupled [17]. Baroreceptors– sensors located in the aortic arch and the carotid sinuses– are responsible for sensing pressure changes of the arterial blood delivered to the systemic circulation and the brain, and relaying the information to the autonomic centers in the brainstem. The term "baroreflex" refers to the negative feedback mechanisms that counteract a drop (or an increase) in arterial BP with an increase (or a decrease) in heart rate, peripheral resistance and venous tone. These changes in turn are reflected in BP due to the mechanical properties (mainly, the capacitance and resistance) of the aorta and large arteries via the so-called Windkessel effect of the systemic vasculature, and other effects related to the cardiac filling [18].

As an illustrative example, consider the time-series of HR and mean arterial pressure (MAP) from a tilt-table experiment [16] shown in Fig. 1-2. Note that the drop in MAP that follows tilting (due to the transfer of blood towards the lower extremities) is followed by a subsequent increase in HR, due to the baroreflex compensatory mechanisms. Moreover, closer visual inspection reveals that the HR time-series exhibits slower oscillations during the tilting events, which are qualitatively similar to the oscillations in the MAP time-series. The tilt-table test has been used by researchers to study orthostatic stress (postural hypotension) [17, 19] and by clinicians to diagnose syncope (or fainting due to low BP) [20]. Beat-to-beat models of HR and BP interactions have been used by researchers to quantify the baroreflex sensitivity and its modifications in response to drugs, and as a consequence of diseases and aging [21].

1.1.3 Clinical Decision Support and Mortality Prediction

Modern clinical data acquisition systems are capable of continuously monitoring and storing measurements of vital signs, such as HR, BP, and respiration (sampled once per second or once per minute), over multiple days of patient hospital stay. As noted previously, the time-series of vital signs exhibit rich dynamical patterns of interaction in response to external perturbations (e.g., administration of drugs or mechanical ventilation) and with onset of disease and pathological states (e.g., sepsis and hypotension). Thus, one would expect to observe a diverse set of dynamical patterns in time-series of vital signs, as the patients go through their course of treatment. However, the large volume and the multidimensional nature of these time-series complicate the task of information integration and assessment of patient progress by clinicians. The current patient acuity scores

(e.g., simplified acute physiology score or SAPS) only include single snap-shot measurements of the vital signs.

Therefore, it is desirable to automatically identify the set of dynamical rules (or simply "dynamics") that govern the interactions among the vital signs within a patient cohort. Each of these dynamics may be categorized as "healthy" or "unhealthy", and may reflect a common phenotypic response to interventions. As such, the discovered dynamics may be useful for prognosis of patient outcome (survival vs. mortality), or for predicting the onset of pathological states. Moreover, if the identified dynamics are constrained to include physiological models of the underlying systems, they may provide clinicians with a mechanistic explanation of the corresponding dynamical patterns, and therefore may suggest personalized treatments to restore healthy dynamics.

1.2 Dynamical Systems in Physiology

In all three examples above, modeling efforts start by identifying some changing properties or attributes of the system that are (1) deemed "important" to the underlying physics, or physiology, of a phenomenon of interest (e.g., periodic breathing), and/or (2) are reflective of the system's response to external (e.g., tilting and drop in BP) or internal perturbations (e.g., onset of sepsis and hypotension). Next, one either acquires measurements of those attributes (or system "variables") or infers their values, and tracks them through time. For instance, in the case of the respiratory control system previously discussed, a researcher may obtain measurements of volume of inspired air, as well as the concentration of CO_2 and O_2 in the inspired and expired air (see Fig. 1-1). Depending on the characteristic time-scale of the underlying mechanisms contributing to the phenomenon of interest, the researcher may extract auxiliary system attributes at coarser levels; for example, via averaging the volume of air over a breath cycle to obtain cycle-averaged time-series of ventilation (also known as minute ventilation).

A mechanistic understanding of the system requires understanding the patterns of interaction among the system variables, and often implies the ability to make predictions about their future values. Therefore, the task of learning in system identification is to reveal the time-dependent rules or dynamics that describe how the future state of a system evolves from its current values, where by "state" we mean a set of (observed or latent) attributes of the system that summarize all we need to know about the system to predict its evolution through time. Then, the problem of inference in system identification is to estimate the system state, given the system dynamics and

observations/measurements of the system variables.

1.2.1 Time-series Modeling in Patient Cohorts

While the examples discussed in the previous sections illustrate the role that the dynamics of interactions among the vital signs play in maintaining a healthy cardiovascular function, more subtle changes to the healthy dynamics of the vital signs may act as an early sign of adverse cardiovascular outcomes. For instance, traditional and nonlinear indices of HR variability (i.e., beat-to-beat fluctuations in heart rate) have been shown to be associated with mortality after myocardial infarction in large cohort studies (several hundred to thousands of patients) [22]. However, these studies fall short of assessing the multivariate dynamics of the vital signs, and do not yield any mechanistic hypothesis for the observed deteriorations of normal variability (that is, they are solely phenomenological). This is partially due to the inherent difficulty of parameter estimation in physiological time-series, where one is confronted by nonlinearities (including rapid regime changes), and measurement artifacts and/or missing data, which are particularly prominent in ambulatory recordings (due to patient movements) and bedside monitoring (due to equipment malfunction). Therefore, in preparation for system identification and parameter fitting, the current paradigm is to isolate stationary segments of the time-series, either by using heuristic rules or statistical tests of stationarity. Next, the researcher may (1) assume simple linear relationships among the system variables and use well-established tools from linear system theory for system identification and parameter fitting [23–26], or (2 use phenomenological/descriptive indices of complexity (e.g., entropy) to quantify the effects of interventions [27, 28]. The linear techniques commonly used have the advantage of providing a more mechanistic description of the system (e.g., using heart rate and blood pressure time-series to infer the baroreflex sensitivity), while the nonlinear descriptors are capable of capturing a richer set of dynamical behaviors, but often lack physiological interpretability in terms of specific underlying mechanisms.

While traditional system identification techniques typically rely on a system's response to external excitations for parameter estimation, external interventions can often alter a subject's physiological state and inadvertently affect the outcomes of physiological experiments. However, due to a large body of accumulated physiological research literature, system identification in physiology has the advantage of having access to informative priors. For instance, when modeling the chemoreflex system, we know that the arterial O_2 and ventilation are tightly coupled (an increase in ventilation results in an increase in arterial O_2 and a decrease in arterial O_2 is typically followed by an increase

in ventilation). Moreover, similar but inverse relationships exist between arterial CO_2 and ventilation. As a result, arterial O_2 and CO_2 are highly anti-correlated, but due to indirect causes (they're both influenced by ventilation). In physiological system identification, a researcher may include such prior knowledge of physiology in the system identification and parameter estimation process, by fixing the coefficients of interactions among arterial O_2 and CO_2 at zeros (or within the Bayesian framework the researcher may place a prior distribution that is highly peaked around zero on the CO_2-O_2 interaction terms).

1.2.2 Switching Dynamical Systems and Cohort Time-series

A central aim of the current work is to develop a framework for automated discovery of evolving dynamics in multivariate physiological time-series from large patient cohorts, such as the Multiparameter Intelligent Monitoring for the Intensive Care II (MIMIC II) database of over 30000 patients [29]. A central premise of our approach is that even within heterogeneous cohorts (with respect to demographic and genetic factors) there are common phenotypic dynamics, corresponding to "healthy" and "unhealthy" physiological states, that influence patient outcomes. Therefore, we define two patients to be "similar" if the multivariate time-series of their vital signs exhibit similar dynamics as the patients go through the course of their treatments, and are subjected to external and internal perturbations (e.g., drugs, or sepsis-induced hypotension).

Specifically, we propose a framework based on the switching linear dynamical systems (SLDS) literature, which allows for piecewise-linear approximation of nonlinear dynamics, as the underlying systems traverse through the various regions within their parameter space. These regions may correspond to different equilibrium points within the phase-space of the underlying dynamical systems, and/or may reflect a common phenotypic response to interventions. Importantly, the framework allows for incorporation of physiologically constrained linear models (e.g., via linearization of the nonlinear dynamics around equilibrium points of interest) to derive mechanistic explanations of the observed dynamical patterns, for instance, in terms of directional influences among the interacting variables (e.g., baroreflex gain or chemoreflex sensitivity), and their individual contributions to the observed system oscillations and/or instability.

1.3 Document Outline and Thesis Contributions

In this thesis, I present a series of techniques for identification and analysis of physiological control systems:

- **Chapter 2** presents a linear closed-loop technique for identification and analysis of physiological control systems. The formulation includes multivariate autoregressive modeling and its extensions to the state-space framework. The techniques developed in this Chapter are useful for stability analysis and identification of sources of oscillations in complex physiological systems involving multiple feedback loops. An extension of the technique to the case of nonstationary systems is also discussed. A signal quality-based technique for robust parameter estimation in the presence of recording artifacts is presented.

- In **Chapter 3** we investigate the feasibility of using fluctuations in multivariate physiological time-series to characterize (oscillatory) behaviors of the underlying systems about their resting stationary equilibrium points. Furthermore, we explore the utility of such characterization to quantify system properties such as stability and the propensity to exhibit oscillatory outputs in response to external disturbances. Using experimental data from an anesthetized, upper-airway bypassed animal model of apnea, we develop a multivariate autoregressive model of the respiratory chemoreflex control system, and use the identification techniques described in Chapter 2 to infer the values of a number of latent factors contributing to breathing instability. Furthermore, we show that together these factors are predictive of the animal's propensity for unstable breathing (or periodic respiration) when subjected to experimental manipulations.

  We then extend our chemoreflex characterization technique to the study of human subjects in a clinical sleep study, where the presence of rapidly changing physiological states that occur during sleep (e.g., sleep-wake transitions, arousal from sleep, changes in controller features with sleeping position, etc.) requires modeling of nonstationary dynamics. We develop an adaptive chemoreflex identification technique that incorporates measures of the quality of experimentally recorded signals into the parameter estimation step, thus mitigating the influence of recording artifacts on the estimated model parameters. We validate our modified chemoreflex identification technique in 21 human subjects, by demonstrating that the technique can detect an increase in ventilatory controller gain produced via an experimental protocol; involving a specialized form of ventilator support. The potential clinical significance of this

work includes the ability to assess respiratory instability in patients with sleep disordered breathing (e.g., Cheyne-Stokes breathing in congestive heart failure, obstructive sleep apnea in adults, and neonatal apnea), and evaluation of weaning from mechanical ventilation in critically ill patients.

- With the goal of modeling dynamical regime transitions in physiological time-series within a patient cohort, **Chapter 4** introduces the switching linear dynamical systems (SLDS) framework. We then present an extension of the framework to incorporate physiological models of the underlying systems, and to discover similar dynamical patterns across a patient cohort. We validate our algorithms using human laboratory recording of HR and BP from a tilt-table experiment, where multiple dynamical patterns are known to exist within the individual subject time-series, as well as across the entire cohort (due to similar maneuvers involving postural changes). Our results demonstrate that the SLDS technique can correctly detect the switching dynamics in time-series of HR and BP, corresponding to various postural changes. Next, we apply the technique to a subset of patients from the MIMIC II intensive care unit database (including $\sim$480 patients with minute-by-minute invasive BP measurements), with the goal of predicting patient outcome variables of interest. We demonstrate that the evolving dynamics of time-series contain information pertaining to survival and mortality of patients, both in hospital as well as up to 30 days after hospital release.

- **Chapter 5** describes an advanced machine-learning technique for outcome-discriminative learning of dynamics within a patient cohort. The main idea of our approach is to present the learning algorithm with the outcomes (or labels) corresponding to each time-series (e.g., survived vs. expired), and to learn time-series dynamics that are most relevant to the discriminative task of distinguishing among the labels. Using the proposed algorithm we demonstrate a significant improvement in decoding postural changes involved in the tilt-table experiment, using the multivariate switching dynamics of HR and BP time-series.

Chapter 2

Closed-loop Identification and Analysis of Physiological Control Systems

2.1 Introduction

A central purpose of mathematical modeling in physiology is to link measurements of physiological variables to the underlying mechanisms. This involves estimating the *parameters* or the *state* of the system from multichannel recordings of the system variables. Here we provide an overview of the class of multivariate autoregressive (MVAR) and state-space linear models for closed-loop identification and analysis of physiological control systems. We present a set of tools for stability analysis, as well as for identification of sources of instability and oscillations in physiological systems involving multiple feedback loops. More specifically, the analysis tools will allow us to (1) determine the characteristics of the individual pathways (or directional links) connecting two or more physiological variables, and (2) assess the contributions of the individual system variables and links to an observed phenomenon of interest (e.g., periodic oscillations within a certain frequency band). Finally, we discuss a nonstationary extension of the MVAR models, and propose a robust parameter estimation technique that weights each incoming measurement in proportion to a measure of signal quality.

2.2 Multivariate Autoregressive Modeling

An MVAR model with maximal lag of P that describes the interactions among K physiological variables can be represented by the following matrix equation:

$$y_t = \sum_{p=1}^{P} \mathbf{a}(p) y_{t-p} + w_t \ , \tag{2.1}$$

where y_t is a column vector of size K×1, the K×K matrices $\mathbf{a}(p)$ contain the autoregressive coefficients, and w_t is a K×1 vector of unexplained residuals, generally defined to have zero mean and covariance Σ. The integer subscript t is usually a time index, but could be (and in our work typically is) a cycle index, e.g., for a cardiac or respiratory cycle. Models of this type have been successfully used by researchers to model the cardiovascular [24,30,31] and respiratory [32,33] control systems, with appropriate physiological constraints imposed on the model coefficients.

Example 1. A trivariate model that describes the interactions among three modeled ventilatory measurements $\dot{V}$, $PCO2$ and $PO2$, which we shall discuss in more detail later in connection with Fig. 3-1, is given by the following specifications:

$$y_t = \begin{pmatrix} \dot{V} \\ PCO2 \\ PO2 \end{pmatrix}_t , \ \mathbf{a}(p) = \begin{pmatrix} a_{\dot{V},\dot{V}}(p) & a_{\dot{V},PCO2}(p) & a_{\dot{V},PO2}(p) \\ a_{PCO2,\dot{V}}(p) & a_{PCO2,PCO2}(p) & 0 \\ a_{PO2,\dot{V}}(p) & 0 & a_{PO2,PO2}(p) \end{pmatrix} , \ w_t = \begin{pmatrix} w_{\dot{V}} \\ w_{PCO2} \\ w_{PO2} \end{pmatrix}_t .$$

The index t here denotes the t-th respiratory cycle, and the indicated variables are the time averages over the corresponding breaths. The matrices $\mathbf{a}(p)$ for $p = 1 \cdots P$ represent the static gains that relate y_{t-p} to y_t; w_t represents the variations in $\dot{V}$, $PCO2$ and $PO2$ that are not explained by the chemical control system properties, and are therefore considered to be the result of external stochastic disturbances to the system, i.e. noise. With these definitions, Eq. (2.1) states that the values of $\dot{V}$, $PCO2$ and $PO2$ at any respiratory cycle t are linear functions of their values at P previous cycles, plus an independent random term w. Each variable is therefore broken down into 4 components, incorporating the additive influence of its own history, the histories of the other two variables, and noise. We assume that the residual vector w is uncorrelated (i.e., is so-called white noise) across different breaths. We also assume the individual elements are uncorrelated with each other and at each t, have mean value zero and constant variances $\sigma^2_{\dot{V}}$, σ^2_{PCO2}, and σ^2_{PO2}, respectively. Thus, σ is diagonal (although this assumption can be relaxed). (In Chapter 3 we use this model to understand the mechanism of ventilatory instability in a lamb model of periodic breathing.)

2.2.1 Parameter Estimation

Least-square-error parameter estimation techniques for the class of MVAR models are well established [34], and are implemented in computational packages such as $Matlab^{TM}$ (see Matlab function *arx.m*). A comparison of the various parameter estimation techniques is presented in *Schlögl* (2006) [35]. Moreover, constrained parameter estimation in Matlab is handled by simply passing a list of fixed model coefficients to the *arx* function.

2.2.2 Transfer Path Functions

Using an identified MVAR model, one may seek to characterize the directional pathways connecting two variables. Equation (2.1) is a discrete convolution; in the frequency domain, this becomes simply multiplicative. Define the Fourier transform of a time signal $q(k)$ by $Q(f) = \sum_k q(k)\exp(-2\pi\sqrt{-1}fk)$; here k indexes the time domain, and f indexes the frequency domain. Our convention is to denote time domain variables in lower case, and transformed variables in upper case. Under the assumption that the signals involved are Fourier Transformable, the Fourier transform of Eq. ((2.1)) is given by:

$$Y(f) = A(f)Y(f) + W(f). \tag{2.2}$$

where from now on the dependence on frequency f will be omitted for notational simplicity. Note that the summation over time in defining $A(f) = \sum_{p=1}^{P} \mathbf{a}(p)\exp(-2\pi\sqrt{-1}fp)$ includes only the first P points, where P is the maximal lag or memory in the model of Eq. (2.1). The i-th row, j-th column entry of A is denoted by $A_{i,j}$.

We now can write an expression for the individual components of the Y in Eq. (2.2):

$$Y_i = [\sum_{j \neq i} T_{j\to i} Y_j] + T_{i\to i} W_i, \tag{2.3}$$

where $T_{j\to i}$, the transfer path (TP) function from the j-th signal to the i-th signal, is defined as

$$T_{j\to i} = \begin{cases} \frac{A_{ij}}{1-A_{ii}}, & i \neq j; \\ \frac{1}{1-A_{ii}}, & i = j. \end{cases} \tag{2.4}$$

The TP functions have physical units; for instance, if j has units of $mmHg$ and i has units of $Lmin^{-1}$ then $T_{j\to i}$ has the units of $Lmin^{-1}(mmHg)^{-1}$ (which is the unit used to quantify chemosensitivity).

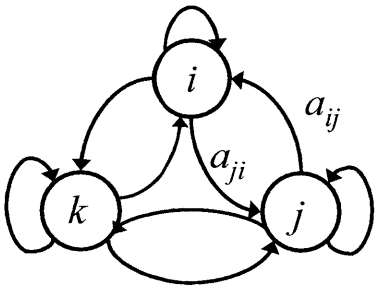

Figure 2-1: A fully connected three-node network.

A physiological system (e.g., the chemoreflex feedback loop) can often conveniently be represented as a network, where the *nodes* represent the physiological variables (e.g., ventilation, arterial CO_2 and O_2), and the connecting *links* represent the interactions between the variables (for instance, the link connecting CO_2 to ventilation may represent the sensitivity of the chemosensors; see also Fig. 3-1).

2.2.3 Fluctuation Transfer Functions

Solving for Y in Eq. (2.2) yields:

$$Y = (I-A)^{-1}W \equiv HW, \tag{2.5}$$

where $H = (I-A)^{-1}$ is a K×K matrix of functions relating the fluctuations in the measured variables to the sources of fluctuation, as defined in Eq. (2.1). Note that the individual components of H can be written in terms of the TP functions defined in Eq. (2.4). For instance, for the three node network shown in Fig. 2-1 we have:

$$H_{j\to i} \triangleq H_{i,j} = \begin{cases} \frac{T_{j\to j}T_{j\to i}+T_{j\to j}T_{j\to k}T_{k\to i}}{1-LG_1^{(1)}-LG_1^{(2)}-LG_1^{(3)}-LG_2^{(1)}-LG_2^{(2)}}, & i\neq j; \\ \frac{T_{i\to i}}{1-LG_1^{(1)}-LG_1^{(2)}-LG_1^{(3)}-LG_2^{(1)}-LG_2^{(2)}}, & i=j, \end{cases} \tag{2.6}$$

where $LG_1^{(1)} = T_{j\to i}T_{i\to j}$, $LG_1^{(2)} = T_{j\to k}T_{k\to j}$, $LG_1^{(3)} = T_{i\to k}T_{k\to i}$, $LG_2^{(1)} = T_{j\to i}T_{i\to k}T_{k\to j}$, and $LG_2^{(2)} = T_{i\to j}T_{j\to k}T_{k\to i}$. Note that Mason's Rule [36] generalizes this expression to an arbitrary number of nodes.

Example 2. For the chemoreflex example above, the components of H pertinent to the influence of

external disturbances on $\dot{V}$ are given by

$$H_{W_{\dot{V}}\to\dot{V}} = T_{\dot{V}\to\dot{V}}/(1-LG)\ [(Lmin^{-1})/(Lmin^{-1})] \quad (2.7)$$

$$H_{W_{PCO2}\to\dot{V}} = (T_{PCO2\to PCO2}T_{PCO2\to\dot{V}})/(1-LG)\ [Lmin^{-1}mmHg^{-1}] \quad (2.8)$$

$$H_{W_{PO2}\to\dot{V}} = (T_{PO2\to PO2}T_{PO2\to\dot{V}})/(1-LG)\ [Lmin^{-1}mmHg^{-1}] \quad (2.9)$$

where

$$LG = T_{\dot{V}\to PCO2}T_{PCO2\to\dot{V}} + T_{\dot{V}\to PO2}T_{PO2\to\dot{V}} = LG_{CO2} + LG_{O2} \quad (2.10)$$

Note that, this simplification is a result of the assumed absence of interaction between $PCO2$ and $PO2$ ($a_{PCO2,PO2} = a_{PO2,PCO2} = 0$), which implies $T_{PCO2\to PO2} = T_{PO2\to PCO2} = 0$.

These equations show clearly that a loop gain near 1 amplifies the influence of the noise terms $w_{\dot{V}}$, w_{PCO2} and w_{PO2}.

A measure of the maximum possible amplification achievable by any combination of external disturbances entering the chemoreflex feedback loop is given by the *matrix 2-norm* of the H matrix (denoted by $||H(f)||_2$). It is frequency-dependent, which implies that the system may selectively amplify certain frequencies in the input disturbances. The location of these frequencies is determined by the *poles* of the H matrix which determine the *natural frequencies* of the system. For the system under study the poles are given simply by the roots of the denominator 1-LG in Eq. (2.6). In general, the system poles are complex numbers of the form $r\ exp(i\theta)$, with magnitude $r \geq 0$ and angle θ. The system is stable if and only if all $r < 1$, i.e., if and only if all poles have magnitude < 1.

2.2.4 Parametric Power Spectra

The fluctuation transfer functions are the building blocks that characterize the contribution of different sources of fluctuation to the power spectra of the individual signals. The power spectrum of the vector time-series y, for the case of white noise residual and covariance Σ at each t, is given in the frequency domain by [37]:

$$S = H\Sigma H^{\intercal}, \quad (2.11)$$

where $^{\intercal}$ denotes matrix or vector transpose.

Example 3. Under the assumption that the fluctuations are uncorrelated across the variables, the covariance matrix Σ has zero off-diagonal elements and diagonal elements $\{\sigma_1^2, \sigma_2^2, \sigma_3^2\} = \{\sigma_{\dot{V}}^2, \sigma_{PCO2}^2, \sigma_{PO2}^2\}$. The diagonal elements of S give the power spectra of the corresponding variables. For instance, the power spectrum of the fluctuations in ventilation is given by

$$S_{\dot{V},\dot{V}} = |H_{W_{\dot{V}} \to \dot{V}}|^2 \sigma_{\dot{V}}^2 + |H_{W_{PCO2} \to \dot{V}}|^2 \sigma_{PCO2}^2 + |H_{W_{PCO2} \to \dot{V}}|^2 \sigma_{PO2}^2. \tag{2.12}$$

This is a weighted combination of the powers due to each individual source of fluctuation, with weights given by the squared magnitudes of the fluctuation transfer functions. Similarly, the off-diagonal elements of the power spectral matrix S give the cross-power term between each pair of variables in the $\dot{V}$, PCO_2 and PO_2 triple. Other derived quantities such as *coherence* [37] can be similarly calculated from the entries of the S matrix.

2.3 Linear Dynamical Systems and State-Space Representation

State-space representations provide a compact and powerful framework for system modeling and analyses; in particular, for characterizing the system response to perturbations in model parameters, for instance as a result of drug administration or physiological state changes. The state-space approach provides a convenient framework for it allows us to quantify the influence of such parameter changes on oscillations of interest and the overall system instability. Moreover, the framework has the advantage of allowing us to conveniently model measurement noise and recording artifacts.

Every MVAR model can be described in a state-space form [38]. For example, one possible state-space representation of Eq. (2.1) is given by:

$$x_{t+1} = Ax_t + e_t \,, \tag{2.13}$$

$$y_t = Cx_t + v_t \,, \tag{2.14}$$

where the state variables x, state noise term e, transition matrix A and observation matrix C are

defined as:

$$x_t = \begin{pmatrix} y_{t-1} \\ y_{t-2} \\ \vdots \\ y_{t-P} \end{pmatrix}, \; e_t = \begin{pmatrix} w_t \\ 0 \\ \vdots \\ 0 \end{pmatrix}, \; A = \begin{pmatrix} \mathbf{a}(1) & \mathbf{a}(2) & \dots & \mathbf{a}(P) \\ I & 0 & \dots & 0 \\ \vdots & \ddots & \vdots & \vdots \\ 0 & \dots & I & 0 \end{pmatrix}, \; C = \begin{pmatrix} I \\ 0 \\ \vdots \\ 0 \end{pmatrix}^{\mathsf{T}}.$$

and v_t represents a measurement noise process that is usually not present in the MVAR model. Under the assumption that w_t and v_t are (white) Gaussian processes we have $e \sim N(0, Q)$, with the upper-right corner elements of Q equal to Σ and zero elsewhere, and $v_t \sim N(0, R)$. Eq. (2.13) represents a linear dynamical system (LDS). More generally, the state-space representation may include nonlinear models of dynamics, and/or observation, as well as non-Gaussian state and measurement noise [39].

2.3.1 Inference in Linear Dynamical Systems

The objective of inference is typically to make optimal estimates (in the minimum mean square error sense) of the state of the system at every point in time, which requires determining state conditional mean $\mu_{t|t}$ and conditional covariance $V_{t|t}$ (Note, the subscript notation $t|t$ is a short-hand for the estimate at time t, given all the observations up to and including time t). The solution to this problem of inference is provided by the celebrated Kalman filter when only data up to current time are used for inference, and the Rauch-Tung-Striebel algorithm (RTS; also known as the Kalman smoother) when all the data are used for inference [40,41].

Appendix A (Section 2.5.1) provides a short derivation of the Kalman filter algorithm. In operator notation the Kalman filter can be represented as follows [41,42]:

$$(\mu_{t|t},\, V_{t|t},\, V_{t,t-1|t}) = KalmanFilter(\mu_{t-1|t-1},\, V_{t-1|t-1},\, y_t;\, A,\, C,\, Q,\, R); \tag{2.15}$$

for t=1 $\cdots$ T. Therefore, given the previous state mean and covariance (with initial values $\mu_{0|0}$, $V_{0|0}$), a new measurement (y_t), and the model parameters ($\{A,\, C,\, Q,\, R\}$), the Kalman filter operator produces the optimal estimates, namely the conditional state mean, covariance, and (optionally) the cross-covariance of the current and the previous state (the latter is used in the backward recursion and by the parameter estimation procedure discussed in the next section). Similarly, the backward recursion step of the Kalman smoother algorithm (or simply the Kalman smoother) in operator

notation is given by [41,42]:

$$(\mu_{t|T},\, V_{t|T},\, V_{t+1,t|T}) = KalmanSmoother(\mu_{t+1|T},\, V_{t+1|T},\, \mu_{t|t},\, V_{t|t},\, V_{t+1|t+1},\, V_{t+1,t|t+1};\, A,\, Q); \tag{2.16}$$

for $t = T-1 \cdots 1$, and with the initial values $(\mu_{T|T},\, V_{T|T},\, V_{T,T-1|T})$ given by the Kalman filter Eq. (2.15).

2.3.2 Parameter Estimation

There are several classes of algorithms available for estimation of the parameters of an LDS $(\{A,\, C,\, Q,\, R,\, \mu_{0|0},\, V_{0|0}\})$, including the Stochastic Subspace Realization algorithm [43,44] and the expectation-maximization (EM) algorithm [41,45]. In this work, we focus on the EM algorithm, since it can be easily extended to the class of switching linear dynamical systems that we will discuss later in this chapter.

EM is a two-pass iterative algorithm: (1) in the expectation (E) step we obtain the expected values of the latent variables (or the state variables in the case of a LDS), and (2) in the maximization (M) step we solve for the set of model parameters that maximizes the complete data log-likelihood. The M-step for a LDS involves solving a series of least-squares problems involving matrices A, C, Q, R. Note that, if the matrix of dynamics A is physiologically constrained, we will need to solve a constrained least squares optimization problem, as discussed in the Appendix (Section 2.5.3). Iteration through several cycles of the EM algorithm is guaranteed to yield the unique maximum likelihood model parameters [45]. A Matlab implementation of the EM algorithm for parameter learning in LDS is available at:
`http://www.cs.ubc.ca/~murphyk/Software/Kalman/kalman.html`.

2.3.3 Selective Modal Analysis

When the physiological system under study can be approximated using a LDS model, and the system is not externally driven, the technique of *selective modal analysis* (SMA) [46] can be used to assess the relative contributions of the individual state variables to the observed system oscillations. In practice, SMA exploits the information embedded in the eigenvectors and eigenvalues of the system to untangle the respective contributions of the various state variables to each oscillatory mode of the system.

From dynamical systems theory we know that the eigenvalues of the A matrix provide us with the complete set of modes (or oscillatory frequencies) of the system. We shall assume the J eigenvalues are distinct. It is straightforward to show that, for a given eigenvalue, the absolute value of the component-wise product of its left and right eigenvectors quantify the relative contributions of each state variable to the corresponding system mode [47]. Let l^j (row vector) and r^j (column vector) be, respectively, the j-th left and right eigenvectors of A, corresponding to the eigenvalue λ_j, defined by:

$$Ar^j = \lambda_j r^j, \tag{2.17}$$

$$l^j A = \lambda_j l^j, \tag{2.18}$$

for $j = 1, 2, \cdots, J$. Note that the eigenvectors are such that $l^i\ r^j = 0$ if $i \neq j$, and we shall assume they are normalized such that $l^j\ r^j = 1$. These eigenvectors allow us to write the undriven solution to Eq. (2.13) as:

$$x_t = A^t x_0 = \sum_{j=1}^{J} \lambda_j^t (l^j x_0) r^j. \tag{2.19}$$

(Note, t indexes discrete time steps, and in modal form we have $x_t = A^t x_0 = (\sum_{j=1}^{J} \lambda_j^t r^j l^j) x_0$). The *Generalized Participation Factor* matrix (**P**) for the j-th mode is defined as [46,47]:

$$\mathbf{P}^j_{k,k'} \equiv r^j_k l^j_{k'}. \tag{2.20}$$

We note the following interpretation of the participation factors:

- If $x_0 = r^j$ (i.e., starting in the subspace corresponding to the j-th mode) then from Eq. (2.19) we have

$$x_t = \sum_{j'=1}^{J} \lambda_{j'}^t (l^{j'} r^j) r^{j'} = \lambda_j^t (l^j r^j) r^j = (\sum_{k=1}^{J} \mathbf{P}^j_{k,k}) \lambda_j^t r^j. \tag{2.21}$$

  Therefore, $\mathbf{P}^j_{k,k}$ measures the relative participation of the k-th state available in building the time-response of the j-th mode. A large $\mathbf{P}^j_{k,k}$ value indicates that the j-th mode is very sensitive to local feedback around the k-th state variable.

- The rate of change of any given eigenvalue λ_j with respect to the a_{pq} element of the A matrix

is given by

$$\frac{\partial \lambda_j}{\partial a_{pq}} = r_q^j l_p^j = \mathbf{P}_{q,p}^j. \tag{2.22}$$

(for a proof see Appendix (Section 2.5.4).) Therefore, $\mathbf{P}_{q,p}^j$ is a measure of the sensitivity of the j-th system mode to variations in the a_{pq} element of the A matrix.

2.4 Modeling Nonstationary Dynamics and Measurement Artifacts

The modeling and analysis tools discussed so far assume that the underlying system parameters and/or the characteristics of the state noise and the measurement noise do not change over time. Next we discuss a nonstationary (or adaptive) extension of the (constrained) MVAR model discussed in Section 2.2.

2.4.1 Kalman Filtering for Modeling Nonstationary MVAR Processes

To accommodate nonstationarity, we may let the autoregressive coefficients vary gradually over time by imposing a random walk on the AR model coefficients [48]:

$$\underline{\mathbf{a}}_{t+1} = \underline{\mathbf{a}}_t + d_t \tag{2.23}$$

where $\underline{\mathbf{a}}$ is a column vector of vectorized AR model coefficients: $\underline{\mathbf{a}} = vec(\mathbf{a}(1), \cdots, \mathbf{a}(P))$. Then the observation equation will take the following form:

$$y_t = H_t \underline{\mathbf{a}}_t + w_t, \tag{2.24}$$

where y_t is a $K \times 1$ vector of modeled variables as defined in Eq. (2.1), H_t is a matrix of the previous values of y $(y_{t-1}, \ldots, y_{t-P})$ structured such that the multiplication of H_t by $\underline{\mathbf{a}}_t$ is equivalent to the summation term in Eq. (2.1). The white-noise term d_t is zero mean with covariance D_t, and controls the degree of smoothness imposed on the $\underline{\mathbf{a}}_{t+1}$ coefficients. The AR noise term w_t is also zero mean with covariance Σ_t. Together Eqs. (2.23), (2.24) define a time-varying MVAR model (in state-space form), and therefore the Kalman filter and smoother algorithms discussed in Section 2.3.1 can be used to solve for the AR coefficients $\underline{\mathbf{a}}_t$, $t = 1, \cdots, T$

Note that in practice the covariance matrices Σ_t and D_t have to be set either *a priori* by the user

or must be learned from the data. Assuming stationary noise, Cassidy and Penny [49] used the EM algorithm discussed in Section 2.3.2 to learn the covariance matrices. In the case of non-stationary noise, various heuristics have been proposed by researchers to update the covariance matrices (see for instance [50]).

2.4.2 Signal Quality Indices

Note that, in contrast to the state-space formulation of the static MVAR process in Eq. (2.13), the state-space formulation in Eqs. (2.23), (2.24) does not include separate noise terms for the autoregressive noise and the measurement noise. Therefore, this formulation cannot distinguish between nonstationary dynamics and nonstationarities that may occur due to measurement artifacts. We propose a technique here to explicitly include information about the quality of measurements within the inference step of the Kalman filter.

With a slight abuse of notation let $\hat{\mathbf{a}}_t$ denote the expected value of the AR coefficients given the measurement stream $y_1, \cdots, y_t$. The Kalman filter estimate $\hat{\mathbf{a}}_t$ is given by (see Eq. (2.32) in Appendix 2.5.1):

$$\hat{\mathbf{a}}_t = \hat{\mathbf{a}}_{t-1} + G_t r_t \,, \tag{2.25}$$

where G_t is a weighting factor (also known as, Kalman gain) which is inversely related to the AR noise covariance matrix Σ (note, the dependence on time is dropped, since for simplicity we assume stationary noise), and $r_t = y_t - H_t \hat{\mathbf{a}}_{t-1}$ is known as the prediction error. Intuitively, Eq. (2.25) states that our best estimate $\hat{\mathbf{a}}_t$ is a weighted combination of our previous estimate $\hat{\mathbf{a}}_{t-1}$ and a second term related to the discrepancy between our model prediction and the actual measurement.

Let us assume that we are given an index of quality for each measurement y_t, denoted by SQI_t (SQI stands for *signal quality index*), which is a number between between zero (very poor measurement quality) and one (excellent signal quality). We make the Kalman gain G_t to be directly (although nonlinearly) proportional to SQI_t, by modifying the Σ as follows [51]:

$$\Sigma'_t = \Sigma \exp(1/SQI_t^2 - 1), \tag{2.26}$$

Σ'_t is a monotonically decreasing function of the SQI_t. The nonlinear weighting function that multiplies Σ approaches one for good quality measurements, at which point the modified covariance matrix Σ'_t will be equal to Σ_t. Note that, if separate signal quality measures are available for each of the K measurements then the diagonal elements of the Σ_t can be individually modified. For the sake

of completeness, we summarize the modified Kalman smoother algorithm:

Initialization

Initialize $\hat{\mathbf{a}}_{0|0}$ and $V_{1|0}$, Σ, D and

Forward-Pass (Kalman Filter)

1. For t from 1 to T

2. Construct the data history matrix H_t

3. Compute: one step prediction of observation: $\widehat{y}_t = H_t \hat{\mathbf{a}}_{t-1|t-1}$

4. Compute residuals: $r_t = y_t - \widehat{y}_t$

5. Update Σ'_t according to Eq. (2.26)

6. Compute Kalman gain: $G_t = V_{t|t-1} H_t^{\mathsf{T}} / (H_t V_{t|t-1} H_t^{\mathsf{T}} + \Sigma'_t)$

7. Compute state estimate: $\hat{\mathbf{a}}_{t|t} = \hat{\mathbf{a}}_{t-1|t-1} + G_t r_t$

8. Compute estimated state covariance: $V_{t|t} = (I - K_t H_t) V_{t|t-1}$

10. Calculate *a-priori* estimated state covariance: $V_{t+1|t} = V_{t|t} + D_t$

END

Backward-Pass (Kalman Smoother)

For t from T-1 to 1

Let $J_t = V_{t|t} / V_{t+1|t}$

1. Compute smoothed estimated state covariance:

$V_{t|T} = V_{t|t} + J_t (V_{t+1|T} - V_{t+1|t}) J_t^{\mathsf{T}}$

2. Compute smoothed state estimate: $\hat{\mathbf{a}}_{t|T} = \hat{\mathbf{a}}_{t|t} + J_t (\hat{\mathbf{a}}_{t+1|T} - \hat{\mathbf{a}}_{t+1|t})$

END

Note, the initial smoothed estimates at step T for the backward-pass are the final state estimates for the forward-pass.

As noted earlier, the covariance matrices Σ and D can be learned from data using the EM algorithm [49]. Given the state sequence from the inference step, the EM solution for Σ and D involves solving two least squares problems [49]. However, inclusion of the SQIs slightly modifies the solution for Σ, since now less importance must be placed on noisy measurements, yielding a weighted least squares problem, with the weight matrix taking the form of a $T \times T$ diagonal matrix with the diagonal elements given by the SQI_i.

One may show that the Kalman estimator effectively weights the past measurements in an exponentially decreasing manner (see Appendix, Section 2.5.2), with a rate inversely related to the

covariance matrix D. Therefore, matrix D controls the degree of smoothness of the sequence of estimated parameters. As such, if we have *a priori* knowledge of the time-scale of model parameter variations, this information can be incorporated into the estimation procedure. In practice, this is done by putting a *Wishart* prior on D [45,49].

2.5 Appendix

2.5.1 Derivation of the Kalman Filter

The Kalman filter provides recursive estimates of the state of a LDS as given by Eqs. (2.13), 2.14). A simple derivation of the Kalman filter (under the assumption of Gaussian system variables) is given by Brown and Hwang (1992) [40], and is summarized here. Our objective is to find the optimal estimate of the state, given all the measurements up to and including time t; we denote this estimate by $\hat{x}_t|y_t, y_{t-1}, \cdots, y_0$. With a slight abuse of notation let x_t denote the state x conditioned on the measurement stream $y_{t-1}, y_{t-2}, \cdots, y_0$. Then the problem of state estimation is equivalent to finding the probability distribution function (pdf) $P(x_t|y_t)$. From Bayes rule we have:

$$P(x_t|y_t) = \frac{P(y_t|x_t)P(x_t)}{P(y_t)} . \tag{2.27}$$

Therefore, the task of estimating the state of the system at time t is reduced to finding the three pdfs on the right-hand-side of Eq. (2.27). The Kalman filter starts by assuming we have an optimal prior estimate, given all the observations up to time $t-1$:

$$P(x_t) \sim N(\mu_{t|t-1}, V_{t|t-1}) \tag{2.28}$$

It follows from Eq. (2.13) that $\mu_{t|t-1} = A_t\mu_{t-1|t-1}$ and $V_{t|t-1} = A_t V_{t-1|t-1}A_t^{\mathsf{T}} + Q_t$. Next, given Eq. (2.14) we arrive at the following pdf for y_t:

$$P(y_t) \sim N(C_t\mu_{t|t-1}, C_t V_{t|t-1}C_t^{\mathsf{T}} + R_t). \tag{2.29}$$

Finally, $P(y_t|x_{t|t-1})$ follows from the fact that if $x_{t|t-1}$ is given, then the conditional pdf of y_t is given by

$$P(y_t|x_t) \sim N(C_t\mu_{t|t-1}, R_t). \tag{2.30}$$

Substituting Eqs. (2.28,2.29, and 2.31) into Eq. (2.27) yields the desired posterior distribution:

$$P(x_t|y_t) \sim \frac{N(C_t\mu_{t|t-1}, R_t)N(\mu_{t|t-1}, V_{t|t-1})}{N(C_t\mu_{t|t-1}, C_tV_{t|t-1}C_t^{\mathsf{T}} + R_t)}. \tag{2.31}$$

This pdf has the mean and covariance:

$$\begin{aligned} \mu_{t|t} &= \mu_{t|t-1} + G_t(y_t - C_t\mu_{t|t-1}), & (2.32)\\ V_t &= [(V_{t|t-1})^{-1} + C_t^{\mathsf{T}}R_t^{-1}C_t]^{-1} \\ &= (I - G_tC_t)V_{t|t-1}\ . & (2.33) \end{aligned}$$

where G_t is a weighting factor (also known as, Kalman Gain) and in given by

$$G_t = V_{t|t-1}C_t^T(C_tV_{t|t-1}C_t^T + R)^{-1}. \tag{2.34}$$

Notice that according to Eq. (2.34) the gain factor G_t and the measurement noise covariance R_t are inversely proportional. Thus, for large measurements noise (as R_t goes to infinity) the gain factor decreases (G_t approaches zero). Therefore, in the presence of reliable observations of the state of the system, the optimal estimator blends the new information with the a priori estimate with more weighs given to the observation. Conversely, if credible observations are not available the optimal estimate using all the data up to the current time is accepted as the new estimate of the state of the system.

2.5.2 Notes on Stability and Effective Memory of the Kalman Filter

Assuming that the Kalman filter has reached a constant Gain condition ($G_t \to G$), Eq. (2.32) for the mean of the estimated state can be alternatively expressed as follows:

$$\begin{aligned} \mu_{t|t} &= (A - GCA)\,\mu_{t-1|t-1} + Gy_t, & (2.35)\\ &= \Gamma\mu_{t-1|t-1} + Gy_t, & (2.36)\\ &= \Gamma^2\mu_{t-2|t-2} + \Gamma Gy_{t-1} + Gy_t, & (2.37)\\ &= \cdots & (2.38)\\ &= \Gamma^b\mu_{t-b|t-b} + G\sum_{j=0}^{b-1}\Gamma^j y_{t-j}, & (2.39)\\ &\approx G\sum_{j=0}^{b-1}\Gamma^j y_{t-j}, & (2.40) \end{aligned}$$

where $\Gamma = (A - GCA)$. The approximation of the last step is justified as follows: for the Kalman filter to be stable, the eigenvalues of the Γ matrix (the roots of the characteristic polynomial for the Kalman filter) must all be less than one in absolute value [40]. The closer to zero these eigenvalues are, the less temporal information is used in the prediction process. This is evident from Eq. (2.40) by observing that the older observations are weighted by progressively smaller weights, since, $\|\Gamma^0\| > \|\Gamma^1\| > \ldots > \|\Gamma^{b-1}\|$, where $\|.\|$ refers to a the *spectral norm*.

2.5.3 Constrained Parameter Estimation

The following Matlab code solves a constrained least squares problem of the form $min_X \frac{1}{2}||XA - B||_2{}^2$, where X, A, and B are $K \times K$ matrices, and some of elements of the matrix X are assumed known (typically set to zero). This type of problems are commonly encountered within the Maximization step of the EM algorithm for the class of linear dynamical systems.

```
function X = ClampedLinearSolver(A, B, FixedList, FixedVals)
% ------------------------------------------------------------
% solves min_X 1/2||X A - B||_2^2 subject to fixed entries given by FixedList and FixedVals.
% A and B are square matrices.
% X is assumed of the form (given a 3 x 3 matrix): [x1 x2 x3;x4 x5 x6;x7 x8 x9]
% FixedList: elements of X matrix that are assumed fixed
% FixedVals: values that the fixed elements of X assume.
% ------------------------------------------------------------
BigA = kron(eye(size(A,1)),A');
B= B';
vecB = B(:) - sum(BigA(:,FixedList) .* repmat(FixedVals,size(BigA,1),1),2);
BigA(:,FixedList)= [];
X = ones(size(BigA,1),1);
Xind = 1:size(BigA,1);
Xind(FixedList)= [];
X(Xind) = BigA \ vecB;
X(FixedList)= FixedVals;
X = reshape(X,size(B,1),size(A,1))';
```

2.5.4 Eigenvalue Sensitivity Analysis

Here we provide a proof of the equality relationship in Eq. (2.22). Note from Eq. (2.17) we have $Ar^j = \lambda_j r^j$. Therefore,

$$\frac{\partial A}{\partial a_{pq}} r^j + A \frac{\partial r^j}{\partial a_{pq}} = \frac{\partial \lambda_j}{\partial a_{pq}} r^j + \lambda_j \frac{\partial r^j}{\partial a_{pq}} .$$

Rearranging the terms we get:

$$\frac{\partial A}{\partial a_{pq}} r^j - \frac{\partial \lambda_j}{\partial a_{pq}} r^j = (\lambda_j I - A) \frac{\partial r^j}{\partial a_{pq}} .$$

Pre-multiplying by l^j causes right-hand-side of the above equality to vanish, yielding:

$$\begin{aligned} l^j \frac{\partial \lambda_j}{\partial a_{pq}} r^j &= l^j \frac{\partial A}{\partial a_{pq}} r^j , \\ \frac{\partial \lambda_j}{\partial a_{pq}} l^j r^j &= l^j \frac{\partial A}{\partial a_{pq}} r^j , \\ \frac{\partial \lambda_j}{\partial a_{pq}} &= l_p^j r_q^j = \mathbf{P}_{q,p}^j . \end{aligned}$$

Chapter 3

Closed-loop Identification of the Respiratory Control System

The part I of this chapter uses data from an animal experiment, involving a lamb model of apnea [52]. The part II of this chapter extends the analysis techniques used in part I to the study of human subjects in a clinical sleep study.

3.1 PART I: Modeling Stationary Dynamics in a Lamb Model of Periodic Breathing

Ventilatory control instabilities in the form of *cyclic apnea* or *periodic breathing* are important in a variety of pathological conditions, including Cheyne-Stokes breathing in congestive heart failure [11–14], obstructive sleep apnea in adult humans [53–55], adult breathing at altitude [7–10] and in neonates [3–6]. Although the specific mechanisms underlying each condition may vary, ventilatory instability can result from increases in the ventilatory sensitivity to hypoxia/hypercapnia (controller gain), in the efficiency of gas exchange (plant gain), or in the circulatory delay between lungs and chemoreceptors [8, 12, 13, 15]. The concept of *loop gain* has been increasingly emphasized because it integrates each of these physiological components into a single function of frequency that describes the stability of the respiratory control system (see Fig. 3-1(A)). When there are delays in the feedback loop, a high loop gain generally describes a system that is intrinsically less stable, whereas a low loop gain describes a more stable system.

The role of loop gain in understanding respiratory control (and in particular, its link to instabil-

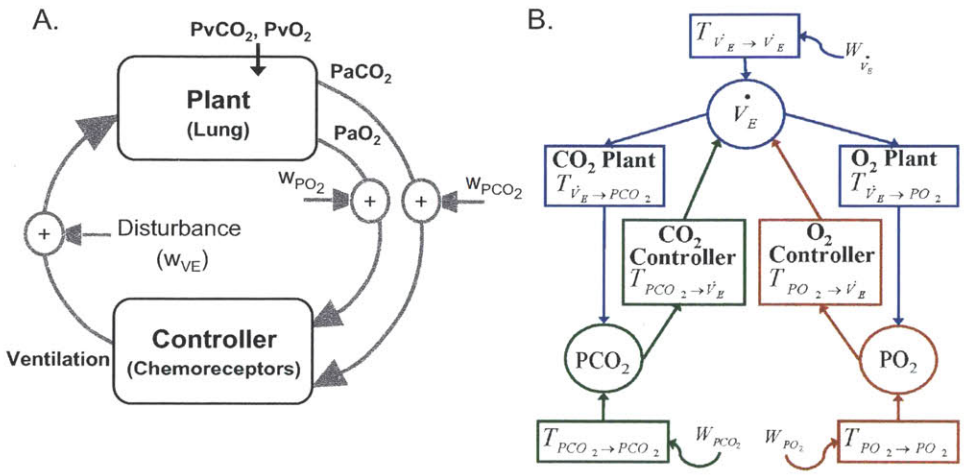

Figure 3-1: Schematic diagram of the closed-loop respiratory control system.
A. The plant represents the gas-exchange dynamics of the lung. Inputs to the plant are the gas tensions in mixed pulmonary capillary blood as well as the level of ventilation, and the outputs are the alveolar/arterial gas tensions. The controller represents the aggregate response of the respiratory pattern generator (RPG) to its inputs (including chemoreceptor outputs, wakefulness drive, etc.). The disturbance term w_{VE} represents random fluctuations in ventilation, for instance as a by-product of noisy neural signal transmission from RPG, variations in upper airway/lung compartment resistance and elastance, variations in the wakefulness drive, and higher cortical inputs [56–58]. The disturbance terms w_{PCO2} and w_{PO2} signify other factors influencing fluctuations in the blood gases, such as changes in metabolic rate, cardiac output, etc. The frequency-dependent product of the plant, the controller gains, and the circulatory delay (not shown here) characteristic is known as the *loop gain* of the system.
B. Closed-loop trivariate model of the ventilatory control system in the frequency domain, modeling the relationships among the measured variables: $\dot{V}$, PCO_2 and PO_2. For instance, $T_{PCO2\rightarrow\dot{V}}$ represents the transfer gain of the pathway involving the CO_2-specific chemoreflexes, responding to fluctuations in P_aCO_2 by a (frequency-dependent) proportional change in minute ventilation; and $T_{\dot{V}\rightarrow PCO2}$ likewise specifies the action of CO_2 exchange in the lung compartment. Similarly, $T_{\dot{V}\rightarrow PO2}$ and $T_{PO2\rightarrow\dot{V}}$ specify the feedforward and feedback components of the O_2-specific loop. All modeled variables are subject to their own source of fluctuation (indicated by the wiggly arrows). Chapter 2 (Section 2.2.2) describes a technique based on autoregressive modeling for characterization of the individual *transfer path functions*.

ity) has been recognized for many years [15, 59]. However, the clinical potential of loop gain has been limited by the fact that it depends upon complex dynamic interactions of central and peripheral chemoreceptor control loops, circulatory delays, and the efficiency of CO_2 and O_2 exchange in the lung. These individual factors are cumbersome and impractical to measure directly [2]. Indirect measures for assessing the risk of periodic breathing include proportional assist ventilation [60], pressure support ventilation [61], and pseudorandom binary stimulation using CO_2 [32]. All existing techniques to characterize the loop gain of the respiratory control system require interventions that alter breathing patterns and blood gases during either quiet wakefulness or sleep. By contrast, in our study the same goal is achieved by taking advantage of spontaneous fluctuations in ventilation $\dot{V}$, and in respiratory gas tensions PCO_2 and PO_2 [62, 63].

The current study, also presented in [33], assessed spontaneous fluctuations in ventilation, end-tidal CO_2 and end-tidal O_2 using a trivariate autoregressive modeling technique, in order to derive estimates of controller and plant gains for the feedback loops controlling both CO_2 and O_2, and thus the overall loop gain of the control system. We tested our method using data from spontaneously breathing anesthetized lambs, before and after administration of domperidone, a drug previously shown to increase the dynamic hypercapnic and hypoxic ventilatory responses, i.e., the associated controller gains [52]. Crucially, if our method has validity, it must yield controller gains that are comparable to published estimates of dynamic hypercapnic and hypoxic ventilatory responses in the lamb. Further, the method would need to reveal an increase in overall loop gain, driven predominantly by an increase in the O_2-specific loop gain, following domperidone administration. In addition to passing these tests, we show that spontaneous fluctuations in ventilation are determined by the loop gain of the respiratory control system [53, 64, 65] and that our analysis closely predicts the cycle-duration of experimentally induced periodic breathing in lambs. Thus the concept of loop gain provides a unifying framework that links the natural variability in ventilation during spontaneous breathing to the cycle-duration of experimentally induced periodic breathing.

3.1.1 Methods and Experimental Setup

Chemoreflex Model Overview

Our model quantifies the inter-relationships among spontaneous fluctuations in breath-to-breath values of ventilation, end-tidal CO_2, and end-tidal O_2 by characterizing the transfer path between each pair of variables (see Fig. 3-1 (B)). The transfer path between O_2 and ventilation defines the con-

troller gain for O_2, and similarly for CO_2. Likewise, the O_2 and CO_2 plant gains are equal to the transfer path functions between ventilation and O_2 and CO_2 respectively. The product of controller and plant gains for O_2, and similarly for CO_2, are the respective loop gains (LG_{O2} and LG_{CO2}). Thus, our model quantifies the two individual feedback loops, one for the feedback control of O_2 and one for control of CO_2. The overall loop gain of the system is equal to the sum of LG_{O2} and LG_{CO2}(as demonstrated in Chapter 2, Section 2.2.3). We note that loop gain is a frequency-dependent quantity [15]. In the current study, where rather than a time index we use a breath index, so duration is represented by the number of breaths, the cycle frequency is described in units of cycles per breath. For example, an oscillation with cycle-duration (or period) of 10 breaths corresponds to frequency of 0.1 cycles/breath.

Using our model, we quantified the impact of the administration of domperidone, a dopamine D_2-receptor antagonist that increases carotid body sensitivity to O_2 and CO_2 [52]. The impact on the control system transfer paths, and thus on loop gain, were assessed using direct measurement of the respiratory variables associated with spontaneous breathing in newborn lambs. Animal data were obtained from previous experiments performed at Monash University, Australia, by Edwards et al. [52]. All surgical and experimental procedures conformed to the guidelines of the National Health and Medical Research Council of Australia and had the approval of the Standing Committee in Ethics in Animal Experimentation of Monash University.

Protocol

The surgical preparation of the animals, and the variables measured and derived in this study, have been described in a previous publication by Edwards et al. [52]. Briefly, 15 newborn lambs (aged 10-20 days) were given a loading dose of ketamine hydrochloride (5mg kg^{-1} Ketamil 100mg ml^{-1}, ILEUM Veterinary Products) via a non-occlusive catheter (Becton and Dickinson, Intracath 19GA) inserted percutaneously into the left jugular vein. A bolus of α-chloralose (80mg kg^{-1}) was then given, followed by a continuous infusion at 20mg $kg^{-1}h^{-1}$. The adequacy of anesthesia was checked by regular stimulation of the inner and outer canthus of the eye and by monitoring heart rate and blood pressure. If needed, supplemental doses of α–chloralose were delivered. Once anesthetized, measurements of respiratory flow, end-tidal P_{CO2} and end-tidal P_{O2} ($PCO2$ and $PO2$) were made. Respiratory flow was measured using a Hans Rudolph (3500A) pneumotachograph. Gas concentrations were measured using a Morgan 901 MK2 CO_2 analyzer and an Ametek S-3A/I O_2 analyzer. Non-occlusive catheters (Datamasters, Victoria, Australia) placed in the right carotid

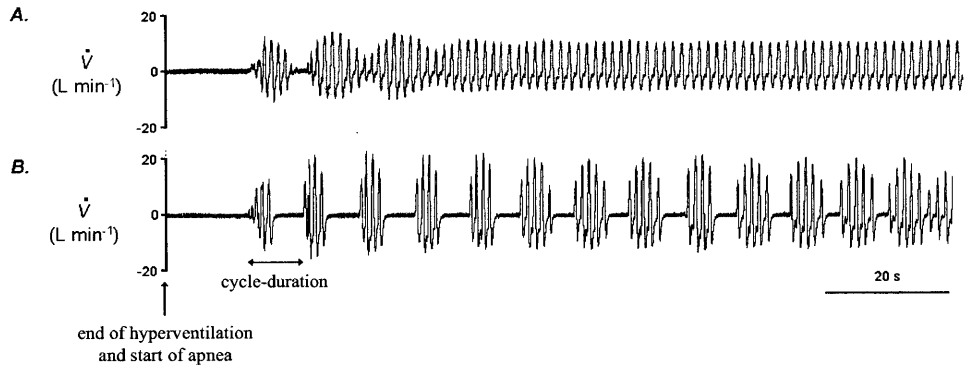

Figure 3-2: Emergence of periodic breathing post-hyperventilation, before (A) and after (B) administration of domperidone.
Cycle-duration was defined as the time between the start of the first breath in a cluster until the start of the first breath in the next cluster, with at least a 5 seconds apnea in between two consecutive clusters [52]. *Epoch-duration* was defined as the time (in seconds) from commencement of breathing following post-hyperventilation apnea to the end of the last apneic pause of periodic breathing. Here, the *number of cycles* of periodic breathing within an epoch under the control (A) and domperidone (B) conditions is 1 and 12, respectively. An increase in the number of cycles of periodic breathing was previously used as a surrogate measure of an increased loop gain [52, 68]. The average cycle-duration (in units of breaths) for each subject (see Table 3.2) was calculated as: epoch-duration (in seconds) / [number of cycles $\times$ average breath duration over the epoch (in seconds/breath)].

artery and jugular vein were used for blood gas sampling. Flow and gas tension signals were acquired at a sampling frequency of 400 Hz. In this study, we used 5-10 minutes of spontaneous breathing data ($\dot{V}$, $PCO2$ and $PO2$) before and at least 10 minute following the intravenous administration of domperidone.

In a previous study Edwards et al. [52] experimentally measured the ventilatory sensitivity to both O_2 and CO_2 (the respective controller gains for O_2 and CO_2) by manipulating inspired O_2 and CO_2 under baseline conditions and again under the influence of domperidone [52]. Additionally, the risk for ventilatory instability under both conditions was characterized using an established model of hyperventilation-induced periodic breathing [66, 67] (see Fig. 3-2). All lambs were sacrificed at the end of the experiment using an overdose of anaesthetic (Lethabarb, Sodium Pentobarbitone, 150mg kg^{-1}; Virbac, Sydney, Australia).

Preprocessing

All data were low-pass filtered at a cut-off frequency of 40 Hz. Respiratory flow recordings were corrected for drift and were integrated breath-by-breath to obtain a continuous tidal volume signal.

The start and end of inspiration were determined by detecting local maxima and minima of the tidal volume signal, and were confirmed visually. Minute ventilation ($\dot{V}$) was calculated for each breath as $\frac{V_T}{T_{tot}}$, where V_T is the tidal volume and T_{tot} is the duration of that breath. To minimize the influence of noise, $PCO2$ and $PO2$ values were calculated by taking the mean of the last quarter of the relevant waveform during expiration . The breath-to-breath time-series data were further high-pass filtered to remove the steady-state baseline and any oscillation in the signals slower than 50 breaths/cycle (using a 7^{th} order Butterworth digital filter with cutoff frequency of 0.02 cycles/breath). The resulting breath-to-breath $\dot{V}$, $PCO2$ and $PO2$ time-series, representing deviations from a steady-state baseline during spontaneous breathing, were used for subsequent analysis.

3.1.2 Trivariate Autoregressive Modeling

Our model represents the three key variables ($\dot{V}$, $PCO2$ and $PO2$), denoting deviations from steady state values, as linear functions of their history (hence the label 'autoregressive'), and random fluctuations. The deterministic component represents the chemical control, including chemoreflexes and gas exchange, while the stochastic component describes how external random fluctuations propagate through the system. The parameters that characterize this system can be utilized to describe the pairwise interactions between the model components [25, 30, 31, 69], and to derive the frequency-domain and stability characteristics of the underlying system. A trivariate model that describes the interactions among the three modeled ventilatory measurements $\dot{V}$, $PCO2$ and $PO2$ can be represented by the following matrix equation (presented earlier in Example 1 of Section 2.2):

$$y_t = \sum_{p=1}^{P} \mathbf{a}(p) y_{t-p} + w_t \ , \tag{3.1}$$

$$y_t = \begin{pmatrix} \dot{V} \\ PCO2 \\ PO2 \end{pmatrix}_t , \ \mathbf{a}(p) = \begin{pmatrix} a_{\dot{V},\dot{V}}(p) & a_{\dot{V},PCO2}(p) & a_{\dot{V},PO2}(p) \\ a_{PCO2,\dot{V}}(p) & a_{PCO2,PCO2}(p) & 0 \\ a_{PO2,\dot{V}}(p) & 0 & a_{PO2,PO2}(p) \end{pmatrix} , \ w_t = \begin{pmatrix} w_{\dot{V}} \\ w_{PCO2} \\ w_{PO2} \end{pmatrix}_t .$$

The vector y_t comprises $\dot{V}$, $PCO2$ and $PO2$ at breath index t; the matrices $\mathbf{a}(p)$ for $p = 1, \cdots, P$ represent the static gains that relate y_{t-p} to y_t; w_t represents the variations in $\dot{V}$, $PCO2$ and $PO2$ that are not explained by the chemical control system properties, and are therefore considered to be the result of external stochastic disturbances to the system, i.e. noise. With these parameters, Eq. (3.1) states that the values of $\dot{V}$, $PCO2$ and $PO2$ at any breath t are a linear functions of their P

previous values, plus an independent random term w. Each variable is therefore broken down into 4 components, incorporating the additive influence of its own history, the histories of the other two variables, and noise. We assume that the individual elements of w are uncorrelated with each other at any breath, and across different breaths; we also assume they have mean value zero and constant variances $\sigma^2_{\dot{V}}$, σ^2_{PCO2}, and σ^2_{PO2}, respectively.

In this work, we fixed the analysis window at 125 breaths and utilized a least-square-error based identification technique [34] to identify the model parameters (see Appendix, Section 3.3.1 for a discussion on model order and window size selection).

3.1.3 Calculation of Controller, Plant, and Loop Gain

The transfer path functions of the control system (controller and plant gains for CO_2 and O_2) describe the direct effect of one variable in the system on another [30]. These functions are derived from the model coefficients a_{ij} (see Chapter 2, Section 2.2.2). Each such function is frequency-dependent (in units of cycles per breath), and is characterized by a magnitude and phase. We denote by $T_{j\to i}(f)$ the frequency-dependent characteristic of the direct pathway connecting the j-th signal to the i-th signal (see Fig. 3-1B). For instance, $|T_{PCO2\to\dot{V}}(f)|$ (units of $\mathrm{Lmin^{-1}mmHg^{-1}}$) describes the magnitude of the change in the component of $\dot{V}$ of frequency f per change in the corresponding frequency component of $PCO2$ (controller gain for CO_2 at frequency f). Similarly, $|T_{\dot{V}\to PCO2}(f)|$ characterizes the magnitude of the change in the component of $PCO2$ at frequency f per change in the corresponding frequency component of $\dot{V}$ (plant gain for CO_2 at frequency f). The CO_2 and O_2 loop gains (denoted by $LG_{CO_2}(f)$ and $LG_{O_2}(f)$, respectively, or with frequency dependence understood) are defined as the product of their respective plant and controller gains (i.e., $LG_{CO_2} = T_{\dot{V}\to PCO2}T_{PCO2\to\dot{V}}$ and $LG_{O_2} = T_{\dot{V}\to PO2}T_{PO2\to\dot{V}}$). Finally, the overall loop gain (LG) turns out to be the sum of CO_2 and O_2 loop gains: $LG = LG_{CO_2} + LG_{O_2}$ (see Chapter 2, Section 2.2.3).

3.1.4 Impact of External Disturbances on Ventilatory Variability: Role of Loop Gain

Our model-based characterization of the ventilatory control system allows for the quantification of the degree to which fluctuations in ventilation result from external fluctuations in $PCO2$, $PO2$, or $\dot{V}$, namely the 'noise' component of these variables that is not explained by the chemical control system (i.e., the w term in Eq. (3.1)). By characterizing the size of the fluctuations in ventilation with respect to the external noise, we can interpret how external disturbances in $PCO2$, $PO2$, or $\dot{V}$ propagate through the feedback loops, resulting in the emergence of oscillatory patterns in

ventilation. This is achieved by using the *fluctuation transfer function* (Chapter 2, Section 2.2.3), denoted by $H_{W_j \to i}(f)$, whereby random disturbances in signal j contribute to the power in signal i. For example, $|H_{W_{PCO2} \to \dot{V}}|^2$ (units of (Lmin^{-1}mmHg^{-1})^2) quantifies the power (i.e., time average of the squared magnitude) of fluctuations in $\dot{V}$ at a given frequency contributed per unit power in the external random disturbance in w_{PCO2} at the same frequency.

The dependence of $H_{W_j \to i}$ on the overall loop gain and the transfer path functions within the various system loops can be described explicitly (see chapter 2, Eq. (2.6)). For instance, we show that the ventilatory response to noise in $PCO2$ is given by $H_{W_{PCO2} \to \dot{V}} = (T_{PCO2 \to PCO2} T_{PCO2 \to \dot{V}})/(1-LG)$. As expected, a smaller noise-buffering effect for $PCO2$ (greater $T_{PCO2 \to PCO2}$) or a greater controller gain for $PCO2$ ($T_{PCO2 \to \dot{V}}$) contributes to a greater impact of noise in $PCO2$ on ventilation. Notably, the finding that $H_{W_{PCO2} \to \dot{V}}$ is dependent on the overall loop gain (LG=LG_{CO2}+LG_{O2}), and specifically on how close the loop gain is to 1, demonstrates that a given extrinsic disturbance in CO_2 results in greater ventilatory variability in a higher gain system, even if overall loop gain is dominated by the O_2 control loop. Thus, if the overall loop gain increases with domperidone, one expects from the analysis that $H_{W_{\dot{V}} \to \dot{V}}$, $H_{W_{PCO2} \to \dot{V}}$, and $H_{W_{PO2} \to \dot{V}}$ will increase in magnitude, and will exhibit a maximal amplitude at the frequency at which periodic breathing is expected to occur (i.e., the *natural frequency* of the system). To obtain the natural frequency of the system, we therefore calculated the frequency at the peak amplitude of the H-matrix within the medium frequency range (see Chapter 2, Section 2.2.3) and compared it to the measured cycle-duration of periodic breathing in individual animals, both before and after domperidone.

3.1.5 Selective Modal Analysis

While the transfer path functions allow us to characterize the directional pathways among two variables, *selective modal analysis* (SMA) quantifies contribution of the individual system variables to an observed system oscillation (see Chapter 2, Section 2.3.3). Here the variables of interest are $\dot{V}$, $PCO2$ and $PO2$ at the different lags. To accomplish this, we convert our trivariate model in Eq. (3.1) to an equivalent state-space form using Eqs. (2.13-2.14), and calculate the participation matrix $\mathbf{P}^j$ for the j-th system mode.

3.1.6 Signal Power as a Measure of Variability

The power spectrum of a signal i (denoted $S_{i,i}$) describes the frequency-dependent variability of the signal, where a prominent peak in the power spectrum indicates the existence of a dominant

oscillatory component. Similarly, the cross-power between two signals i and j (denoted $S_{i,j}$) is a frequency-dependent function that quantifies their cross-correlation properties. From our model, analytical forms for $S_{i,i}$ and $S_{i,j}$ are derived (and given in Chapter 2, Section 2.2.4). An empirical estimate of the power spectrum of a signal, for comparison with the analytical expression, can be obtained by using the periodogram averaging technique, computed via the fast Fourier transform of the signal [65]. A close match between the model-based and the periodogram-based spectra indicates that the model provides an adequate description of the periodic components of the data.

3.1.7 Data Analysis and Statistics

To facilitate comparison between control and domperidone conditions, we considered three distinct frequency bands: low frequency (LF) oscillations with periods 16-50 breaths/cycle, medium frequency (MF) oscillations of 5-15 breaths/cycle, and high frequency (HF) oscillations from 2-4 breaths/cycle. Within each frequency band the magnitude of the derived quantities (controller gains, plant gains, loop gains, $S_{i,j}$, etc.) were averaged, and the resulting indices compared between the control and domperidone studies, using the Wilcoxon signed rank test with significance level $p < 0.05$. Data are presented as medians and interquartile ranges. Since the MF band spans the range of cycle-durations of periodic breathing observed experimentally [32], we focus on the values of controller and plant gain, and the overall loop gain, measured in this band. For the purpose of predicting the predisposition to ventilatory instability, we consider loop gain magnitude in the MF band to be an indicator of nearness to instability, with a value >1 suggesting an unstable system, and a value < 1 suggesting a stable system. This indicator attempts to capture in a simple way, suited to our purpose, the essence of the Nyquist criterion [15].

Results

3.1.8 Respiratory Variables and Experimentally Derived System Properties

The average baseline respiratory variables derived from the breath-by-breath time-series are shown in Table 3.1. The administration of domperidone resulted in a slight increase in $\dot{V}$, mainly due to a decrease in T_i ($p < 0.05$) and a near significant decrease in T_{tot} ($p = 0.09$). These changes were accompanied by a decrease in $PCO2$ ($p < 0.05$), and an increase in $PO2$ ($p < 0.05$).

3.1.9 Trivariate Analysis Results

Controller, Plant, and Loop Gain

Transfer path analysis revealed a significant increase in the magnitude of O_2 chemosensitivity ($|T_{PO2\rightarrow\dot{V}}|$) across the LF and MF bands following administration of domperidone (see Fig. 3-3). Figure 3-4 shows the comparison of the control and domperidone studies for the average gain magnitudes within the MF range for each subject. We found a significant increase in the median value of the average MF band gains in $|T_{PO2\rightarrow\dot{V}}|$. Although the median value of the average MF band gain in $|T_{PCO2\rightarrow\dot{V}}|$was increased with domperidone, this failed to reach statistical significance. The complete analysis results for the CO_2- and O_2-specific loop gains, and the overall LG, are presented in Table 3.2; both the LG_{O2} magnitude and the overall LG magnitude increased significantly in the LF and MF regions. Notably, the LG_{O2} magnitude under both control and domperidone conditions was greater than LG_{CO2} magnitude, and the overall LG was comparable to the LG_{O2}. Moreover, a two-fold increase in the overall LG magnitude with domperidone was observed in conjunction with the greater propensity towards periodic breathing (1/15 vs. 13/15 animals exhibited periodic breathing; see Table 3.2). Fig. 3-6A shows that the model-based CO_2 and O_2 controller sensitivities in the MF band compared favorably with the experimentally measured ventilatory sensitivity to CO_2 and O_2, respectively ($p < 0.05$).

Impact of External Disturbances on Ventilatory Variability: Role of Loop Gain

Figure 3-6 shows the average fluctuation transfer function (H) spectra that quantify the influence of external disturbances on $\dot{V}$. Note the increase in the amplitude of the spectra and emergence of a sharp peak in the medium frequency (MF) band after administration of domperidone. Figure 3-5B is a plot of the cycle-duration of experimentally induced periodic breathing vs. the cycle-duration of these dominant MF band peaks, showing good agreement between the predicted and experimentally measured cycle-durations (predicted: 7.9 [7.0, 9.2], measured: 7.5 [7.2, 9.1] breaths). Moreover, Fig. 3-6C demonstrates a significant increase in the number of experimentally observed periodic breathing cycles with an increase in the overall LG (r^2= 0.50, $p < 0.05$).

Power and Cross-Power

Fig. 3-7 shows the model-based power spectra of the individual signals overlaid on their periodogram-based power spectra. Before the administration of domperidone the $\dot{V}$ spectrum (Fig. 3-7A) shows

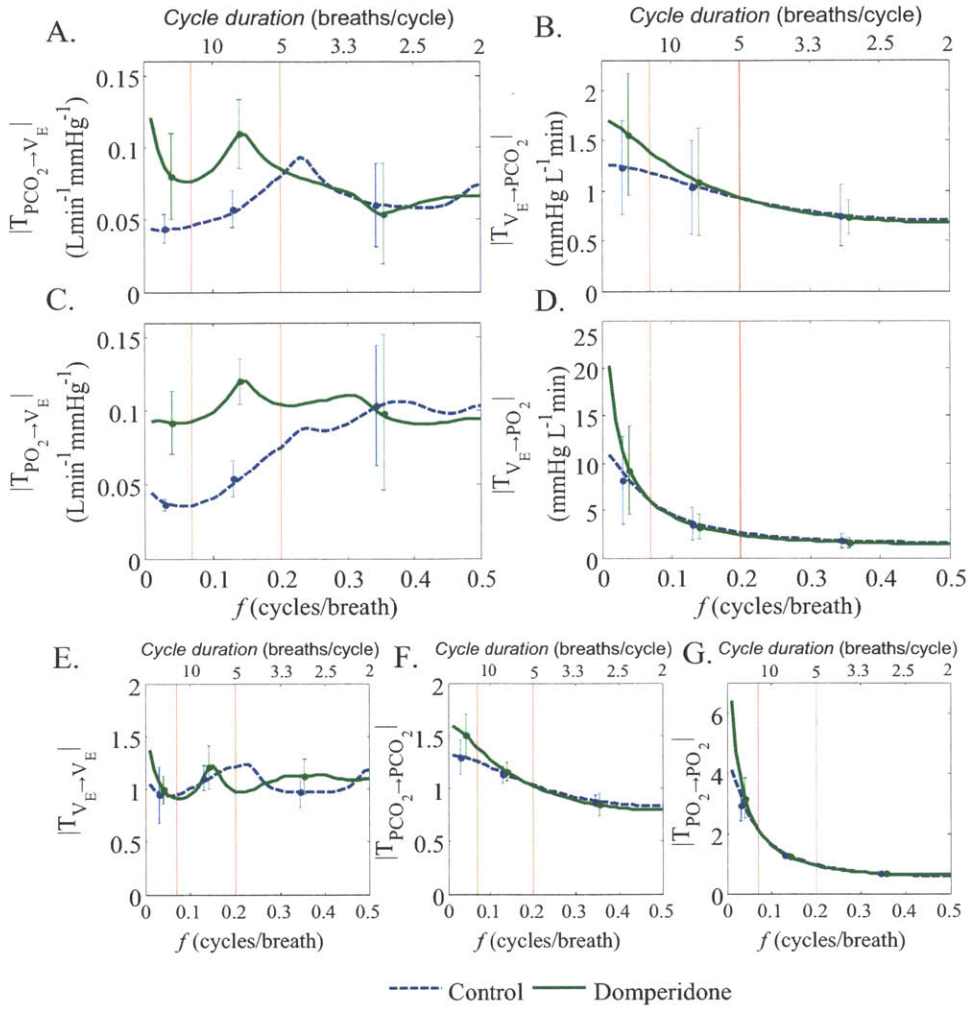

Figure 3-3: Transfer path analysis results

Group averages and standard deviations for the transfer path functions between various components of the system (Panel A: CO_2 controller gain, B: CO_2 plant gain, C: O_2 controller gain, D: O_2 plant gain, E, F, and G: self-influencing transfer paths for $\dot{V}$, $PCO2$ and $PO2$, respectively. Corresponding plots for the control (dashed black lines) and domperidone (solid grey lines) studies are subperimposed on the same panel. In each case, the bottom horizontal axis represents the frequency in units of cycles/breath, the top horizontal axis represents the cycle-duration (or period) in units of breaths/cycle. The boundaries of the LF, MF, and HF bands are marked by the vertical lines. Noteworthy is the marked increase in O_2 controller gain magnitudes ($|T_{PO2\rightarrow\dot{V}}|$) within the MF region after the administration of domperidone. For each subject, the average values of the gain magnitudes for each transfer path function within the MF frequency bands were calculated for assessment of the influence of domperidone (see Fig. 3-4).

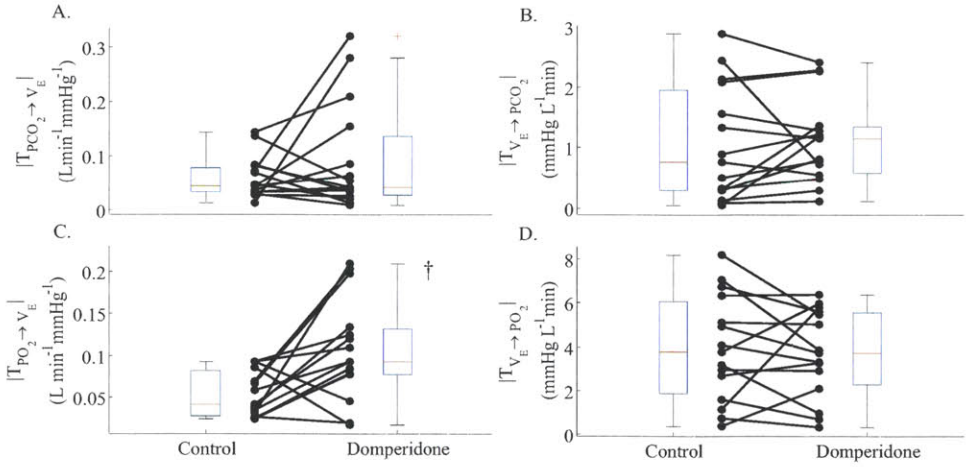

Figure 3-4: Comparison of average transfer path gain magnitudes within the MF band between the control and domperidone studies.
Each black closed circle signifies one of the 15 subjects. Within each panel results from the control and domperidone studies are pooled together and the associated box-plot is displayed next to each group. The group median controller gain for CO_2 (panel A) did not change significantly between control and domperidone studies (0.045 [0.032, 0.079] vs. 0.054 [0.028, 0.137] $Lmin^{-1}mmHg^{-1}$) while the median gain for O_2 (panel C) increased significantly (0.058 [0.032, 0.090] vs. 0.093 [0.077 0.132] $Lmin^{-1}mmHg^{-1}$, $p < 0.05$). Furthermore, the CO_2 plant gain (panel B) changed from 0.76 [0.296, 1.945] to 1.157 [0.590, 1.348] $mmHgL^{-1}min$ and the O_2 plant gain (panel D) changed from 3.759 [1.872, 6.053] to 3.962 [2.273, 5.576] mmHg $L^{-1}min$.

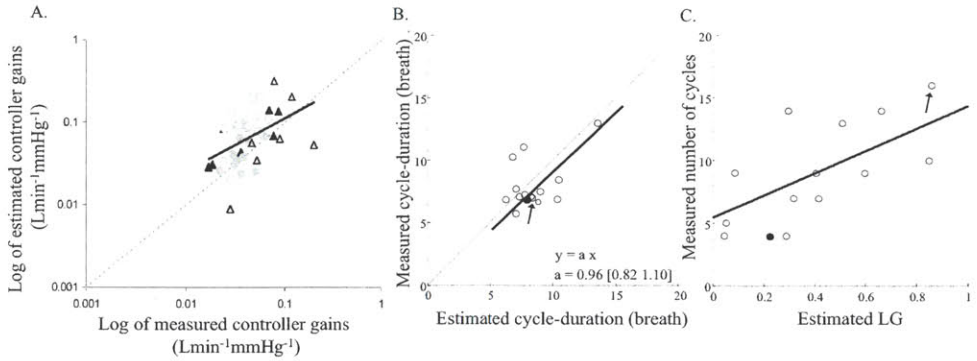

Figure 3-5: Comparison of experimental results and model-based findings using spontaneous breathing.

A. Comparison of the experimentally measured [52] and calculated sensitivities for CO_2 (triangles) and O_2 (circles) in both control (filled) and domperidone conditions (open). The line of identity (dashed) and the regression line (solid) are superimposed on the graph (r^2= 0.43, $p < 0.05$). B. Comparison of the predicted (or estimated) and measured cycle-durations under control (closed circle) and domperidone conditions (open circles). The slope a = 0.96 of the fitted line (heavy black) did not differ significantly from the line of identity (thin grey). C. Plot of the measured number of cycles of periodic breathing vs. the estimated overall LG, indicating a significant increase in the number of periodic breathing cycles with an increase in the overall LG ($r^2 = 0.5$, $p < 0.05$). In panels B and C the arrow indicates post-domperidone cycle-duration and number of cycles of the only subject that exhibited periodic breathing under the control study (see also Table 3.2).

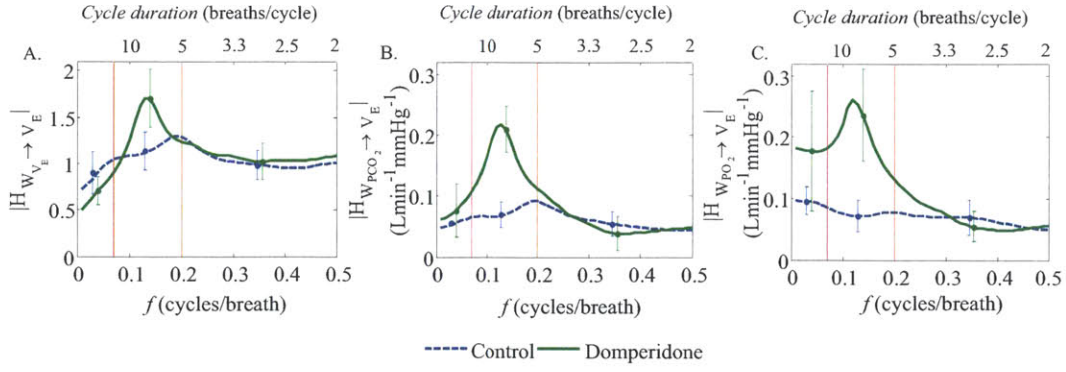

Figure 3-6: fluctuation transfer function.

Group averages and standard deviations for the fluctuation transfer function magnitudes for minute ventilation, for the control (dashed black) and domperidone (solid grey) studies (top three panels). Noteworthy is the strong peak around 8 breaths/cycle after administration of domperidone in panels A, B, and C, indicating that fluctuations in any of the variables may induce periodic variability in minute ventilation, with a cycle-duration of approximately 8 breaths/cycle. The cycle-duration of this medium frequency peak for the individual subjects is compared against the experimentally measured cycle-durations in Fig. 3-5.

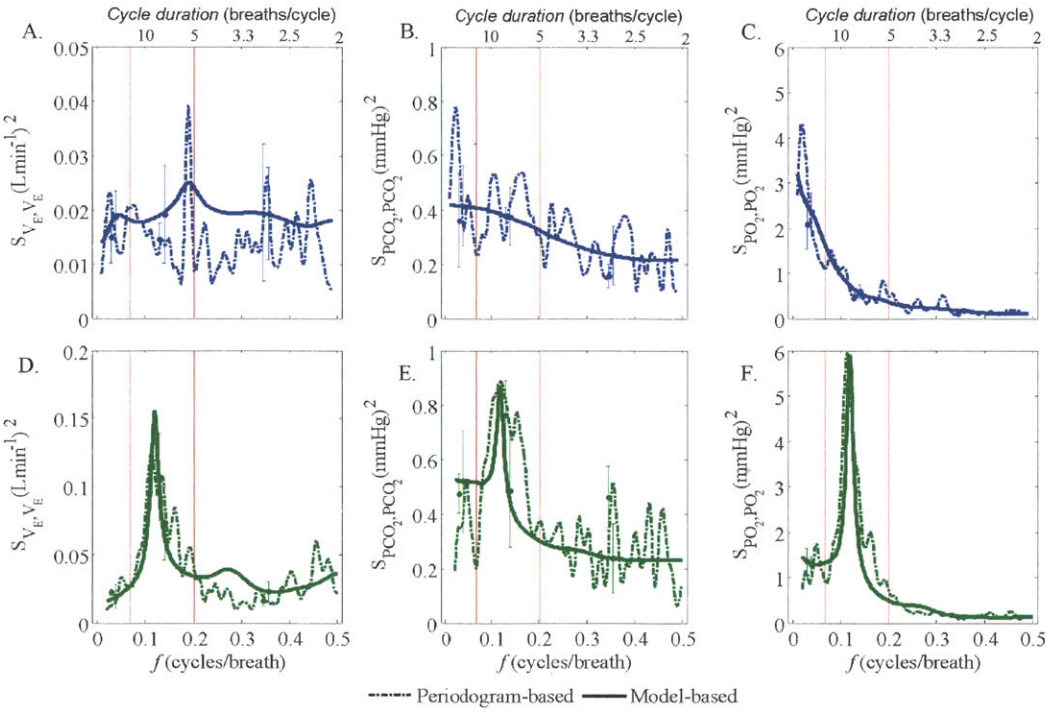

Figure 3-7: Power spectrum.

Group averages and standard deviations for the power spectrum of the individual time-series for the control (black plots in panels A, B and C) and domperidone (grey plots in panels, D, E and F) studies. In each case, the dashed-lines are the periodogram-based power spectra (using a rectangular window of size 100 samples with 25 samples overlap) and the bold solid lines are the model-based power spectra. Note that the frequencies of the dominant peaks are consistent with reports [65] in resting human subjects.

Variable	Control	Domperidone
V_T **(mL)**	71.38 [66.28, 79.69]	74.32 [63.6319, 77.95]
T_I **(sec)**	0.62 [0.57, 0.69]	0.60 [0.54, 0.62] †
T_{TOT} **(sec)**	1.65 [1.49, 1.82]	1.58 [1.38, 1.70]
$\dot{V}$ **(L/min)**	2.61 [2.22, 3.40]	3.13 [2.66, 3.84]
$PCO2$ **(mmHg)**	42.50 [40.00, 46.55]	38.46 [36.29, 39.07] †
$PO2$ **(mmHg)**	99.69 [96.36, 104.60]	103.71 [100.99, 105.47] †

Table 3.1: Baseline variables and experimentally derived system parameters.
Median and first and third quartiles (inside the brackets) of tidal volume (V_T), inspiratory time (T_I), total breath duration (T_{TOT}), minute ventilation (V_T/T_{TOT}), end-tidal carbon dioxide ($PCO2$), and end-tidal oxygen ($PO2$). The † indicates a statistically significant change in the median values between the control and domperidone studies ($p < 0.05$).

a peak around the cycle-duration of 5.5 breaths/cycle, although the signal power is relatively low. After the administration of domperidone (Figs. 3-7D-F) all three signals exhibit marked periodicity around the cycle-duration of 8 breaths/cycle. Moreover, under domperidone (Fig. 3-7D) the model-based and the empirical power spectra of $\dot{V}$ closely match.

Selective Modal Analysis

Figure 3-8 summarizes the application of the SMA technique to the spontaneous breathing time-series recording before and after domperidone administration. Application of the Wilcoxon signed rank test (with false discover rate correction for multiple comparison) revealed a significant increase in contribution of $PO2$ at lags one and two to the MF band natural frequency. This result confirms and explains the experimental observation that only administration of supplemental O_2 (not CO_2) resulted in abolishment of periodic breathing-under the Domperidone protocol- in these lambs.

3.1.10 Discussion

The major finding of our study is that there is sufficient information in the natural fluctuations in $\dot{V}$, $PCO2$, and $PO2$ during spontaneous breathing to 1) estimate, through the use of a physiologically based dynamic model, the magnitude of the dynamic hypercapnic and hypoxic ventilatory responses, and 2) predict the propensity to periodic breathing (associated with high loop gain), and the corresponding cycle-duration of periodic breathing, should it occur. The power of our analytical method is that it permits quantitative determination of controller and plant gains for both CO_2 and O_2, and thus overall loop gain of the respiratory control system, without the need for externally

Subject	Control			Domperidone		
	$\lvert LG_{CO2} \rvert$	$\lvert LG_{O2} \rvert$	$\lvert LG \rvert$	$\lvert LG_{CO2} \rvert$	$\lvert LG_{O2} \rvert$	$\lvert LG \rvert$
1	0.02	0.02	0.02 (-, -)	0.01	0.41	0.40 (6.8, 7.0)
2	0.06	0.12	0.18 (-, -)	0.18	0.31	0.41 (8.4, 7.3)
3	0.00	0.02	0.02 (-, -)	0.01	0.35	0.36 (6.1, 10.4)
4	0.02	0.07	0.06 (-, -)	0.02	0.08	0.09 (6.7, 11.0)
5	0.04	0.16	0.14 (-, -)	0.00	0.05	0.05 (8.8, 8.1)
6	0.18	0.37	0.21 (-, -)	0.09	0.49	0.49 (-, -)
7	0.1	0.29	0.22 (8.1, 7.1)	0.22	0.75	0.84 (8.5, 7.2)
8	0.07	0.17	0.13 (-, -)	0.02	0.3	0.31 (10.5, 8.6)
9	0.01	0.13	0.14 (-, -)	0.00	0.08	0.08 (-, -)
10	0.02	0.24	0.22 (-, -)	0.00	0.1	0.10 (13.4, 13.1)
11	0.01	0.34	0.34 (-, -)	0.09	0.62	0.65 (7.1, 6.4)
12	0.21	0.18	0.37 (-, -)	0.27	0.62	0.59 (10.4, 7.1)
13	0.28	0.57	0.47 (-, -)	0.12	0.77	0.83 (7.5, 7.3)
14	0.06	0.16	0.14 (-, -)	0.38	0.6	0.51 (7.9, 7.5)
15	0.00	0.11	0.11 (-, -)	0.04	0.26	0.30 (7.4, 7.7)
Median Gains	0.04 [0.01, 0.09]	0.16 [0.11, 0.28]	0.14 [0.12, 0.22]	0.04 [0.01, 0.17]	0.35 †† [0.14, 0.62]	0.40†† [0.15, 0.57]
Cycle-durations	(8.1, 7.1)			(7.9 [7.0, 9.2], 7.5 [7.2, 9.1])		

Table 3.2: Subject-by-subject comparison of loop gain magnitudes and cycle-durations. The magnitude of LG_{CO2}, LG_{O2} and overall LG in the MF band, in addition to predicted and experimentally observed cycle-durations (in units of breaths), for the control and domperidone studies. Both the LG_{O2} and the overall LG increased significantly with domperidone in the MF as well as LF band (although, the latter are not presented here). The †† indicates a statistically significant increase in LG with domperidone ($p < 0.005$). For each subject, the first number inside the parenthesis is the predicted cycle-duration of periodic breathing (in units of breaths) and the second number is the experimentally observed cycle-duration (in units of breaths), with the dashed lines (–) indicating that either the subject had no peak in the MF band or that the subject did not exhibit periodic breathing under the hyperventilation protocol. The first and the third quartiles associated with the median values are given inside the brackets.

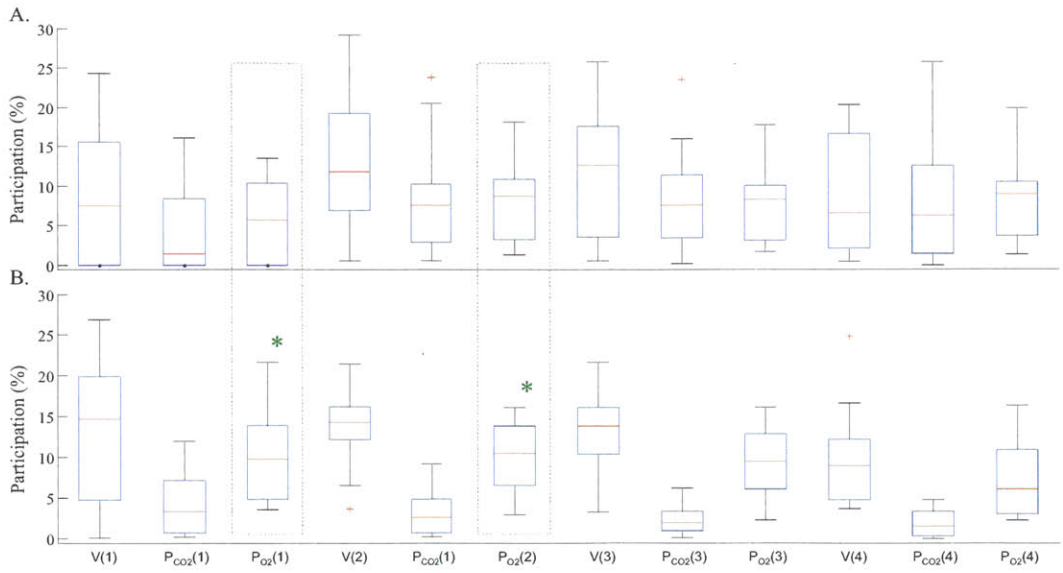

Figure 3-8: Selective modal analysis results.
The first four lags of each of the three ventilatory variables constitute the 12 state-variables shown on the horizontal axis. Panels A and B show the percent participation of each state-variable (pooled across all 15 subjects and presented as a boxplot) to the natural frequency of the system within the medium frequency (MF) band, before and after administration of domperidone.

applied perturbations of the respiratory system. Our results are consistent with earlier findings in newborn lambs, that administration of domperidone increases the loop gain, predominantly by an increase in the O_2 controller gain, making the respiratory control system more prone to instability [52]. We have demonstrated quantitatively that the controller gains for O_2 and CO_2 in the lamb have comparable magnitudes, but a higher plant gain for O_2 results in the overall loop gain being dominated by the O_2 feedback path rather than that for CO_2. Application of selective modal analysis also confirms the major contribution of O_2 to the observed system oscillations. This finding is consistent with experimental observation in the same animals, that stabilizing the O_2 feedback path during periodic breathing, but not the CO_2 feedback path, stabilizes ventilation [70]. By virtue of being based entirely on non-invasive measures of spontaneous breathing, our approach represents a new methodology to 1) assess the risk of ventilatory instability in vulnerable populations, 2) identify the primary causes of instability in individual patients and design appropriate treatment, and 3) follow patient progress in response to therapeutic interventions.

Comparison with other approaches: estimated vs. measured hypoxic/hypercapnic responses. In a previous work, by presenting different concentrations of inhaled CO_2 and O_2, Edwards et al. [52] determined the corresponding controller gains under control conditions and following domperidone

administration. They observed that CO_2 response to three breaths of CO_2 increased with domperidone [52]. However, in our study, we found no statistically significant increase in CO_2 controller gain. This disparity is likely to be due to the large perturbation used previously (14% inspired CO_2) whereas the current method relies on spontaneous oscillations in CO_2 in order to calculate CO_2 controller gain. Nevertheless, the MF band CO_2 and O_2 transfer path gains were significantly correlated with the experimental (traditional) estimates of controller gain (see Fig. 3-6A), indicating that our trivariate analysis sufficiently captured the measured changes. Interestingly, the experimental eupneic controller gain for O_2 was not significantly different between control and domperidone (data not shown), despite the difference previously described at hypoxic levels [52]. Importantly, our novel method was sufficiently sensitive to detect a significant increase in MF band O_2-controller gain during spontaneous breathing with the administration of domperidone (see Fig. 3-4). As such, our method accurately captures changes in sensitivities of the respiratory controller to O_2 from spontaneous breathing records, obviating the need to perform labor-intensive and technically challenging experiments.

Comparison with other approaches: independent laboratories. The link between respiratory variability during spontaneous breathing and the chemoreflex feedback system in adult human subjects has been previously investigated [53, 64]. Of note is the work of Van den Aardweg and Karemaker [65], who observed that mean inspiratory flow time-series in healthy human subjects contains oscillatory components with a cycle-duration of approximately 10 breaths/cycle. Based on a simplified model of the CO_2 chemoreflex feedback loop, they concluded that the correlations in breath-to-breath values of the mean inspiratory flow may reflect the response of the chemoreflex feedback system to (uncorrelated) noisy disturbances in the partial pressure of arterial carbon dioxide (P_{aCO_2}). More recently, approaches using spontaneous changes in ventilation and CO_2 have been developed, with one study demonstrating that opioid administration (remifentanil) substantially increases plant gain but reduces controller gain [71]. These studies assume that fluctuations in O_2 have negligible impact on ventilation, and have not shown a direct link between spontaneous breathing variability and the occurrence or cycle-duration of periodic breathing. Our work goes beyond these previous studies in two important ways. First, the use of a trivariate analysis incorporates the influence of $PO2$ fluctuations as well as $PCO2$ fluctuations on changes in $\dot{V}$. As such, our method has general appliability since it does not make any prior assumptions regarding the dominance of one loop vs. the other. Second, our study goes beyond a phenomenological description of the ventilatory variability and examines the mechanism responsible for the amplitude and frequency characteristics of

the observed ventilatory periodicity.

Our study confirms the view [64,72] that fluctuations in breath-to-breath values of $\dot{V}$ provide insight into ventilatory instability. Thus increasing the gain of the system with domperidone causes an increase in magnitude of the ventilatory variability per unit of extraneous noise in P_{O2}, P_{CO2} and ventilation (as reflected in the fluctuation transfer functions for $\dot{V}$) within the medium frequency band. We also show explicitly that the closer $|LG|$is to 1, the greater the ventilatory variability that will result from any given amplitude of external perturbation. Specifically, small random perturbations in ventilation (i.e. noise) of amplitude A will become approximately augmented by 1/(1-LG) to exhibit fluctuations of amplitude A/(1-LG); for example, if $LG = 0.9$ then an underlying variability with a mean amplitude of 0.5 L/min will be amplified 10 fold to exhibit oscillations of mean amplitude of 5 L/min. Hence, while periodic central apneas will consistently occur in a control system with $|LG| > 1$, in the appropriate frequency band, ventilatory oscillations are also expected to occur in a system with moderately high gain (loop gain less than 1, but close to 1) in the presence of sufficient underlying ventilatory variability. Importantly our data demonstrate (Fig. 3-5) that the frequency content of this ventilatory variability reflects the natural frequency of the system, as confirmed by our finding that cycle-duration predicted from spontaneous breathing corresponds closely with the cycle-duration of experimentally induced periodic breathing. In addition, the experimentally observed number of cycles of periodic breathing increases as the estimated overall $|LG|$ approaches 1. Although a large $|LG|$ is only one of the several factors contributing to an increase in the number of cycles, the number of cycles has been previously used as a surrogate measure of an increased loop gain [68, 73], because the number of cycles is expected to increase as the damping of oscillations decreases, i.e., as $|LG|$ approaches 1. Our results indicate that over 50 percent of the variation in the number of cycles of periodic breathing can be explained by the overall $|LG|$ estimates during spontaneous breathing. Taken together, our evidence leads us to conclude that the periodic components of spontaneous ventilatory variability and periodic breathing are emergent properties of a high gain respiratory control system, amplifying the external disturbances entering the chemoreflex feedback loop.

Methodological considerations

The model presented here is based on a number of simplifying assumptions. Firstly, it is assumed that $PCO2$ and $PO2$ are accurate proxies for their corresponding arterial values [74]. However, arterial and end-tidal $PO2$ can differ substantially when gas exchange is not ideal. On the basis that

$PCO2$ and $PO2$ are strongly correlated with their respective arterial values, our measures of loop gain are not expected to be affected unduly by this assumption. However, a constant ratio of arterial and end-tidal $PO2$, as observed dynamically in the lamb [75], is expected to yield an underestimated controller gain and an overestimated plant gain compared with the true arterial values.

Our second assumption was that the fluctuations in the current values of $\dot{V}$ are linearly dependent on their own history and fluctuations in the previous values of $PCO2$, and $PO2$, and also that the $PCO2$ and $PO2$ fluctuations are only dependent on their own respective histories along with the previous values of $\dot{V}$ fluctuations. Thus, we assumed a negligible synergistic interaction between $PCO2$ and $PO2$ at the controller or the plant, which we considered valid for the case of small perturbations in $PO2$ and $PCO2$ around resting levels. However, such an interaction may be important in the case of large disturbances in the human adult, when a large hypoxic dip would be expected to increase the controller gain for CO_2. For our application such $PO2$-$PCO2$ interaction was deemed negligible, since the controller responses to $PCO2$ and $PO2$ are additive in the lamb [70]. Thirdly, although the calculated values for CO_2 controller gain in the mid-frequency range could be a combination of both peripheral and central chemosensitivity to CO_2, it is likely that the relatively long brain tissue wash-in/wash-out time constants for CO_2 [76–79] limit the central chemoreceptors from contributing significantly to rapid or transient changes resulting from breath-to-breath variability.

From a model-fitting prospective, non-stationarities may arise as a result of nonlinearly interacting processes at time-scales corresponding to fluctuations in cardiac output, respiratory rate, sleep-state or arousal-related changes in respiratory mechanics and the chemical control system, behavioral factors, and those induced by interventions. In the current study, we addressed this issue by using a relatively short analysis window to minimize the impact of non-stationarities. Moreover, a strong collinearity (or high correlation) between modeled variables may be a confounding factor in determining the true contribution of $PCO2$ and $PO2$ to ventilation. In our data, the maximum correlation coefficient between $PCO2$ and $PO2$ (over the range of -4 to 4 breath lags) was -0.48 [-0.38 -0.64] under baseline conditions and -0.61 [-0.41 -0.78] after domperidone administration. This intermediate degree of correlation between breath-by-breath fluctuations in $PCO2$ and $PO2$ is expected from respiratory system physiology, given the different buffering capacities of the body for CO_2 and O_2.

Reliability of the estimated model parameters relies on the existence of adequate perturbations in the frequency bands of interest so that every possible pathway is sufficiently excited. Notably, the spontaneous variability was sufficient in our study to achieve a close correspondence between

the spectral characteristics of our model output and the experimental data (Fig. 3-7), and the correspondence appeared to improve as ventilatory variability was increased following domperidone administration. Under conditions of insufficient natural fluctuations, it may be beneficial to subject the system to external perturbations; in principle, our autoregressive method is also applicable to the pseudorandom binary sequence CO_2 stimulation technique [32].

We explicitly assume that the pathways in Fig. 3-1B are the dominant ones, but are mindful that the respiratory control system could be affected by fluctuations in heart rate, blood pressure, cerebral blood flow, and sympathetic activity. For instance, Henry et al. [80] reported a possible coupling between the chemoreflex and baroreflex mechanism, in which case extra pathways representing the interactions among heart-rate, blood-pressure and respiration need to be incorporated. However, this particular effect is likely to be minimal in the current work, as the lamb data we used came from a study in which mean arterial blood-pressure was stable [52]. Regarding other confounding factors, we simply note that the concordance between our estimates and the experimental measurements (in both the cycle-duration of periodic breathing and in the independent assessments of hypoxic and hypercapnic respiratory responses) strongly suggests that the model we propose captures the essentials of the linear feedback mechanisms of the respiratory control system.

Utility of Our Analytic Approach

Our non-invasive method for quantifying the dynamic hypercapnic and hypoxic ventilatory responses and overall loop gain potentially has widespread clinical applicability. Increased loop gain as a phenotypic trait in obstructive sleep apnea is predictive of the response of patients to treatment with supplemental oxygen [55]. An elevated hypoxic ventilatory response at high altitude is associated with the presence of unstable breathing [7]. Moreover, heightened dynamic hypercapnic responses delineate heart failure patients with periodic breathing in the form of Cheyne-Stokes respiration from those patients that do not exhibit ventilatory instability [2, 81]. Importantly, in heart failure patients, elevated hypoxic and/or hypercapnic sensitivity [82] as well as the presence of Cheyne-Stokes respiration itself [83] are important predictors of mortality. Thus, model-based characterization of the respiratory control system using spontaneous breathing provides a simple non-invasive diagnostic tool, allowing potential therapies to be directed at the dominant feedback loop (i.e. CO_2 vs O_2 loop) responsible for the instability.

Conclusion

This study shows that it is possible to characterize and quantify the dominant feedback loops within the respiratory control system from measurements of fluctuations in ventilation and respiratory gas tensions during spontaneous ventilation. Consequences of this approach include the ability to assess risk of ventilatory instabilities, cycle-duration of periodic breathing, and hypoxic and hypercapnic ventilatory responses. The major advantages of our method are that it requires minimal subject intervention and it makes no demand on subjects to comply with complex breathing protocols. Through overall loop gain, our method promises new ways to assess the potential for respiratory instabilities, especially important in cases of congestive heart failure or obstructive sleep apnea. Importantly, identification of the mechanisms responsible for an increase in loop gain in such individuals would enable clinicians to target particular therapies on an individualized basis, and to monitor patient response to such therapies.

3.2 Part II: Modeling Nonstationary Dynamics in Human Research Polysomnography Recordings

The autoregressive technique discussed in the previous sections fits only a single model to the entire dataset. However, the presence of rapidly changing physiological states that occur during sleep (e.g., sleep-wake transitions, arousal from sleep, changes in controller features with sleeping position and state) requires a continuous adaptive model, as these are the times of greatest clinical and scientific interest. The primary goal of this study is to adaptively identify a model, with an associated adaptive measure of loop gain and its components, in human subjects in a clinical sleep study setting. In the current study, we aimed to validate our modified loop gain measurement technique in human subjects by testing whether it could detect an increase in loop gain produced by proportional assist ventilation (PAV). PAV is a mode of ventilatory support that works by generating pressure at the airway in proportion to a person's instantaneous respiratory effort; more effort yields more support [60]. Thus, PAV amplifies the ventilatory response to a given level of CO_2, and therefore increases controller gain and thereby loop gain (see Fig. 3-9).

To facilitate this study and to extend our technique to human subjects in a clinical setting, we implemented several modifications to our methodology. First, we developed an adaptive loop gain estimation technique that accounts for nonstationarity in chemoreflex system characteristics (e.g., as a result of changes in sleep stages, body position, etc). Second, we proposed a framework for including measures of the quality of the experimentally recorded signals into our adaptive estimation technique, thus mitigating the influence of recording artifacts on the estimated chemoreflex system parameters. We show that with the proposed modifications, our method is able to track chemoreflex system non-stationarity associated with an increase in controller gain via PAV, effectively deal with measurement related artifacts, and provide clinically useful indices of plant, controller, and loop gain.

3.2.1 Experimental Setup and Methods

Thirteen continuous positive air pressure (CPAP)-treated obstructive sleep apnea (OSA) subjects (age: 45 ± 10) were recruited from the sleep laboratory at Brigham and Women's Hospital, Boston. All OSA subjects had a history of OSA with an apnoea/hypopnoea index (AHI) >10 events/hr during supine NREM sleep, and a documented CPAP use of >5 hours per night for at least two months prior to the study. Eight healthy controls (age: 35 ± 10) were also recruited from adver-

tisements. All subjects gave written, informed consent before participation in this study, which was approved by the Human Research Committee of the Brigham and Women's Hospital. Subjects were excluded if they were taking any medication known to influence breathing, sleep/arousal or muscle physiology. Additionally, subjects were excluded if they had a history of renal failure, neuromuscular disease or other major neurological disorders, uncontrolled diabetes, heart failure, central sleep apnea/Cheyne-Stokes respiration, uncontrolled hypertension, thyroid disease, or any other unstable medical condition. Female subjects were screened to ensure that they were not pregnant.

Protocol

Subjects underwent a clinical polysomnography (PSG) to measure OSA severity, and a research PSG to measure LG using PAV. Subjects were asked to sleep supine for the majority of the night with the standard clinical montage that included electroencephalogram (EEG; C3/A2, O2/A1), electrooculography, and submental electromyography. Oxygen saturation was monitored using a pulse-oximeter attached to the index finger (BCI Capnograph Series 9000, Waukesha, WI), and PCO_2 was sampled at the mask by using a calibrated infrared CO_2 analyzer (BCI). Data were collected and stored using the Alice digital PSG system (Philips Respironics, Murrysville, PA). Sleep state, arousals, and respiratory events were scored by a blinded sleep technician according to standard AASM criteria.

For the research PSG, electrodes were attached similarly to the clinical PSG and were fitted with a nasal mask (Gel Mask; Respironics, Murrysville, PA) attached to a pneumotachometer (model 3700A; Hans-Rudolph, Kansas City, MO) for measuring airflow. Airway pressure was measured at the mask using a pressure transducer (Validyne, Northridge, CA). CO_2 was continuously recorded from a catheter placed inside the nostril with a capnograph (Vacumed, Ventura, CA). The mask was connected to a BiPAP Vision mechanical ventilator which was capable of delivering continuous positive airway pressure (CPAP) alone or in combination with proportional assist ventilation (PAV). During supine non-rapid eye movement (NREM) sleep, the tidal-volume amplification factor ($V_T AF$) of PAV was slowly increased until periodic breathing was induced [55]. However, in most individuals, periodic breathing could not be induced despite maximum levels of PAV and often trials were terminated due to occurrence of an arousal.

For the analysis we only considered those portions of the PAV data before periodic breathing or arousal was observed. We estimated loop gain and its components from 5 to 10 minute long segments of spontaneous breathing (CPAP only) and compared them to the values measured from

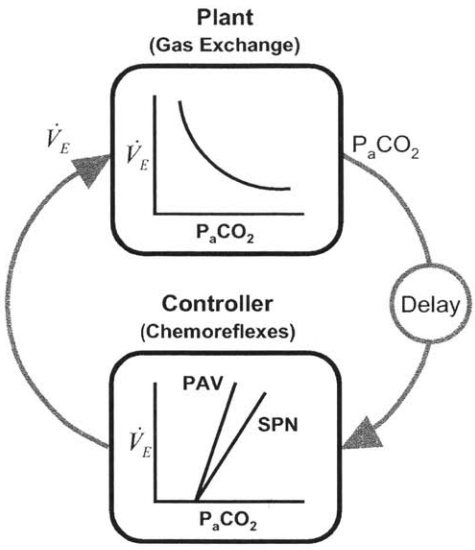

Figure 3-9: Effect of PAV on the chemoreflex feedback loop.
Proportional assist ventilation (PAV) works by generating pressure at the airway in proportion to a person's instantaneous respiratory effort, resulting in an increase in controller gain. The (frequency-dependent) product of the various components around the loop (plant, delay, and controller) is known as the loop gain of the system. A high loop gain causes the system to become unstable, whereas a low loop gain corresponds to a more stable system. The plant characteristic depicted above is known as the "metabolic hyperbola", and represents the steady-state relationship between ventilation (as input) and $PaCO_2$ (as output).

segments of equal duration on PAV, when loop gain is known to be elevated [60].

Signal Preprocessing and Signal Quality Assessment

Signals for flow and P_{CO2} were processed according to the steps described in the previous study to extract the breath-by-breath time-series data for minute ventilation ($\dot{V}$) and end-tidal P_{CO2} ($PCO2$); $\dot{V}$ was calculated for each breath as, $V_T/overT_{tot}$ where V_T is the tidal volume and T_{tot} is the duration of the breath. Because $PCO2$ can sometimes be inaccurate due to low expiratory volume and mask leak (which would confound the estimation of loop gain), a breath-by-breath signal quality index for $PCO2$ (SQI_{CO2}) was created (see Fig. 3-10 panel E and Fig. 3-11), as follows. We fit a line through the end-tidal portion of the $PCO2$ waveform for each breath, and the slope (S) of the line was recorded. A flat line (i.e., S=0) indicates that the measured end-tidal CO_2 is of high quality, while non-zero slopes correspond to end-tidal CO_2 estimates with low signal quality. Next, for every breath, if the corresponding $PCO2$ was different by more than 5 mmHg from the preceding $PCO2$ then a penalty of 0.2 points was imposed; if different by more than 10 mmHg, a penalty of 0.4 points was imposed. The final value of the SQI_{CO2} was set equal to max (0.5, 1-S-penalty). For

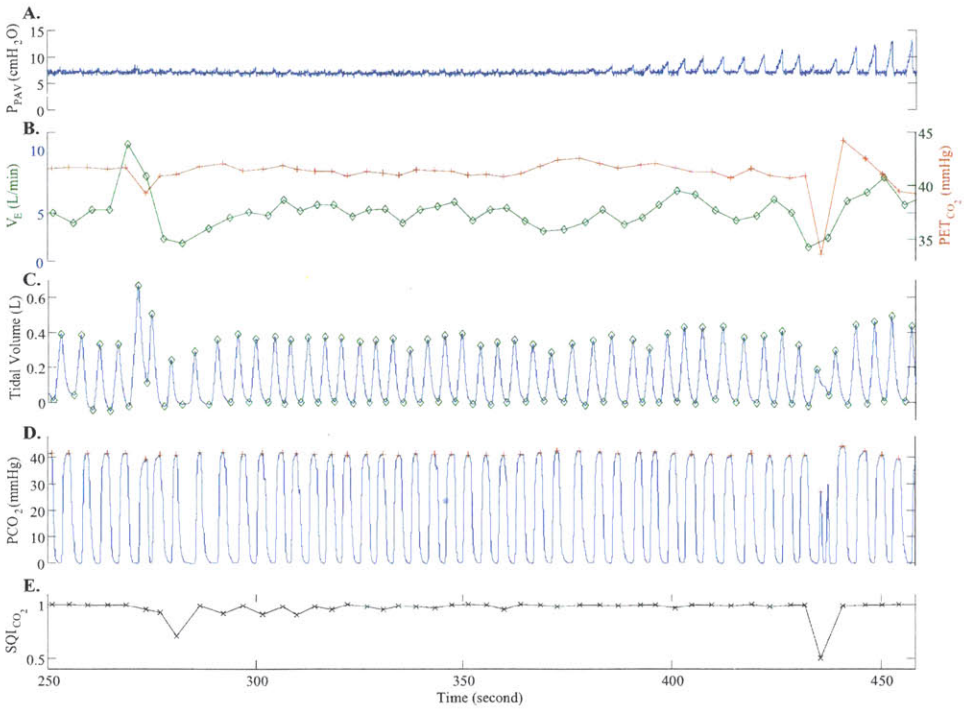

Figure 3-10: Example of recorded waveforms and derived time-series.
A. Pressure setting of PAV. This subject was breathing on 7 cmH_2O CPAP until PAV was initiated (time = 380 s) and slowly ramped up, as evident by the increasingly larger inspiratory pressures delivered. B. Derived minute ventilation ($\dot{V}$; open circles) and end-tidal PCO_2 ($PCO2$; + symbols) time-series. C. Tidal volume waveform with marked breath onset and end-inspiration. D. PCO_2 waveform with marked $PCO2$ values. E. A derived breath-by-breath index of $PCO2$ signal quality (SQI_{CO2}). Note that, due to a shallow breath at ~430 seconds, $PCO2$ drops suddenly to below 35 mmHg. However, since the end-expiratory portion of the waveform is not flat, the corresponding $PCO2$ is a poor estimate of the alveolar PCO_2 level. The associated SQI_{CO2} for this breath is 0.5, indicating that this breath should be trusted less when estimating the chemoreflex system parameters.

example, a slope of S=0.15, and a 6 mmHg change in $PCO2$ (penalty= 0.2), results in an SQI_{CO2} of 1-0.15-0.2 = 0.65.

3.2.2 Adaptive Calculations of Controller, Plant, and Loop Gain

A nonstationary time-series is one whose statistics (e.g., mean, variance) change over time. This is relevant to the current study in which PAV, and therefore loop gain, was slowly increased over time. Furthermore, during spontaneous breathing, non-stationarities may arise as a result of nonlinearly interacting processes at time-scales corresponding to fluctuations in cardiac output, sleep-state, body position, etc. Thus, an adaptive (or time-varying) model would be desirable, not only for the PAV data presented here, but also for the spontaneous breathing as well.

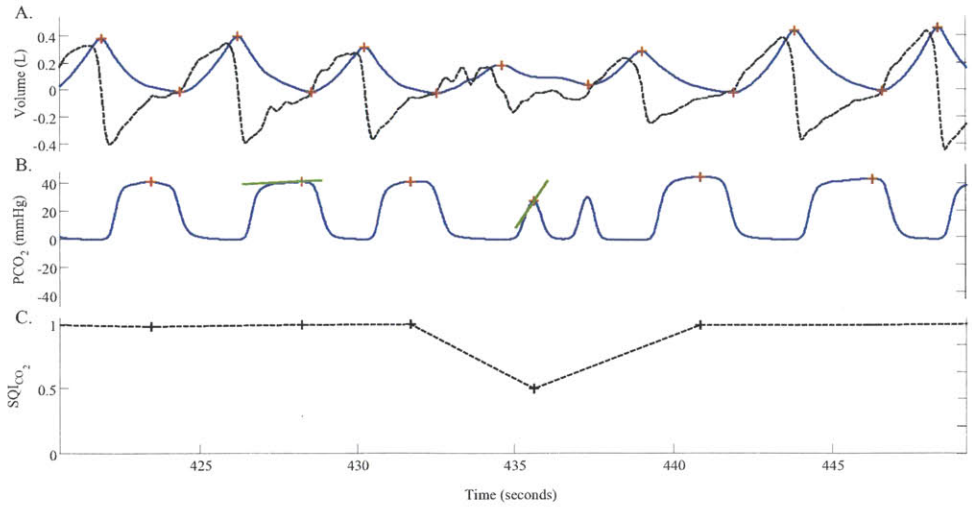

Figure 3-11: Signal quality index for CO_2.
Panel A shows a flow signal (dark, dotted line) and the corresponding tidal-volume signal (blue, solid line) with marked inspiratory onsets and ends. Note that the *PCO2* corresponding to breath #4 is about 28 mmHg. Such a sudden drop in *PaCO2* is physiologically infeasible. As one can see in panel A, the low value of *PCO2* is likely due to an incomplete emptying of the lungs. Panel C shows the calculated signal quality for each breath (+ marks).

To obtain an adaptive estimation of the chemoreflex system parameters, we used the time-varying analog of the constrained chemoreflex model presented in Section 3.1.2, by converting Eq. (3.1) to its equivalent state-space form using the methods discussed in Chapter 2, Section 2.4. In order to minimize the deleterious effects of end-tidal *PCO2* artifacts (which occurs due to shallow or hypopneic breathing), we used our breath-by-breath indices of signal quality for *PCO2* to lower the contribution of poor signal quality breaths in our estimation procedure (as discussed in Chapter 2, Section 2.4.2). Figure 3-12 shows an example PAV trial (panel A), breath-by-breath signal CO_2 signal quality time-series (panel B), and the corresponding adaptive estimates of controller, plant and loop gain, before (panel C) and after (panel D) including signal quality indices.

We initialized the initial state estimates $\underline{\mathbf{a}}_{0|0}$ and $V_{1|0}$, and the autoregressive noise covariance matrix Σ by fitting (static) autoregressive models to the first fifty breaths from each of the subjects and then averaging the resulting parameters to obtain a single set of initial values that were used for all the subjects. Next, we used the EM algorithm to learn the covariance matrices Σ and D from the data, where the latter covariance determines how fast the parameters are adapted, see Section 2.4. Given the identified model parameters, and the smoothed state estimates ($\underline{\mathbf{a}}_{t|T}$, $t = 1, \cdots, T$), for each t we derived estimates of controller, plant, and loop gain using the methods described in Section 3.1.9.

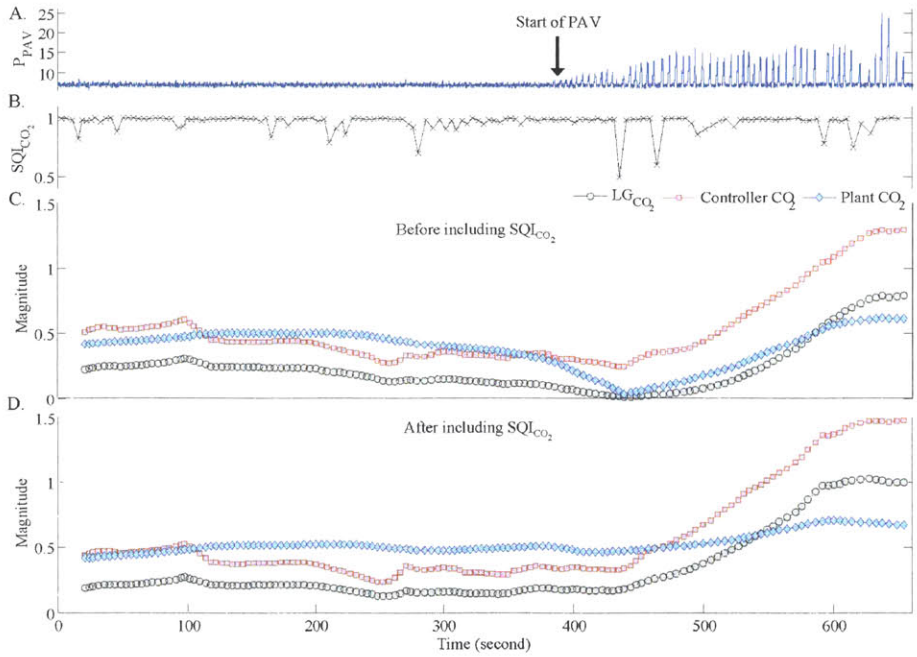

Figure 3-12: Adaptive estimation of controller, plant, and loop gain of the proposed chemoreflex model.
Starting around 400 seconds, the PAV tidal-volume amplification factor (V_TAF) was progressively increased (panel A). Our derived breath-by-breath index of $PCO2$ quality (SQI_{CO2}) is shown in panel B. Adaptive estimation of the controller, plant, and loop gain of the system before and after including the SQI_{CO2} are presented in panels C and D, respectively (the confidence bounds associated with individual estimates are omitted for clarity). The gains are all computed as average over the MF range, as in the previous study. Inclusion of the SQI_{CO2} in panel D resulted in a better estimate of the plant gain and the loop gain around the 430 seconds mark (note that the controller gain drops to zero in panel C due to artifactual measurements of end-tidal CO_2 around the 430 seconds time mark). As expected, application of the PAV resulted in a progressively larger value of the controller gain, which resulted in progressive increase in system loop gain.

Subjects	Baseline			PAV		
	Control (N=8)	OSA (N=13)	Overall (N=21)	Control (N=8)	OSA(N=13)	Overall (N=21)
V_T (L)	0.39 ± 0.07	0.43 ± 0.09	0.42 ± 0.08	0.42 ± 0.08	0.47 ± 0.09	0.45 ± 0.09†
T_I (sec)	1.77 ± 0.20	1.77 ± 0.31	1.77 ± 0.27	1.82 ± 0.28	1.73 ± 0.31	1.76 ± 0.29
T_{TOT} (sec)	4.30 ± 0.53	4.20 ± 0.72	4.24 ± 0.64	4.37 ± 0.65	4.30 ± 0.71	4.33 ± 0.67†
$\dot{V}$ (L/min)	5.43 ± 0.59	6.27 ± 1.03*	5.96 ± 0.96	5.77 ± 0.74	6.71 ± 1.22*	6.35 ± 1.14†
$PCO2$ (mmHg)	39.92 ± 3.77	40.68 ± 2.92	40.39 ± 3.20	38.83 ± 3.56	39.40 ± 3.14	39.18 ± 3.2†

Table 3.3: Break-down of respiratory variables for control and OSA subgroups.
* $p<0.05$ between OSA and Controls † $p<0.05$ between baseline and PAV

Statistical Analysis

Paired t-tests were used to compare the baseline CPAP only against CPAP+PAV on controller, plant and loop gain within individuals and were performed using Matlab. Furthermore, unpaired t-tests were used to compare controller, plant and loop gain between healthy controls and OSA subjects. A p value of less than 0.05 was considered significant. Values are presented as means ± S.D.

Results

3.2.3 Respiratory Variables and Experimentally Derived System Properties

The average respiratory variables during spontaneous breathing (CPAP only) and PAV are shown in Table 3.3. PAV induced a small increase in V_T and T_{TOT} with the net effect of a slight increase in $\dot{V}$ and decrease in $PCO2$. Table 3.3 also shows a break down of the respiratory variables according to the patient population (control vs. OSA). OSA subjects exhibited a significantly higher $\dot{V}$, both during baseline and on PAV.

3.2.4 Effect of PAV on Controller, Plant, and Loop Gain

Figure 3-13 provides a comparison of loop gain (panel A), controller gain (panel B) and plant gain (panel C), for the individual subjects between baseline CPAP and CPAP+PAV. The autoregressive technique revealed a statistically significant increase in loop gain ($p<0.05$), predominantly driven by an increase in controller gain ($p<0.005$). The final value of V_TAF was in general lower in OSA (3.2±0.8) vs. control (3.9±0.9), although the difference did not reach statistical significance (p=0.16).

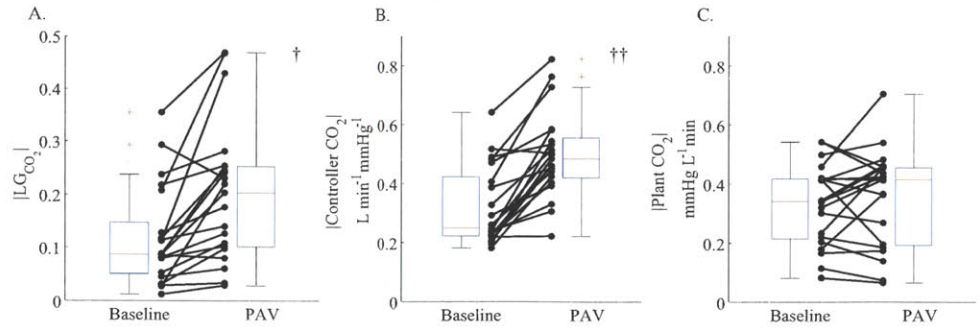

Figure 3-13: Group comparisons of controller, plant, and loop gain (mixed control and OSA subjects).
Baseline CPAP and CPAP+PAV loop gains (panel A), controller gains (panel B), and plant gains (panel C). The group average of the loop gain increased from 0.140.10 to 0.210.13, predominantly due to a significant increase in controller gain from 0.355 ± 0.145 to 0.536 ± 0.145 $Lmin^{-1}mmHg^{-1}$. The group average of the plant gain did not change significantly (from 0.35 ± 0.13 to 0.33 ± 0.14 mmHgL-1min). The † and †† indicate a statistically significant difference ($p<0.05$ and $p<0.005$, respectively).

3.2.5 Baseline Controller, Plant, and Loop Gain in OSA vs. Controls

Compared to controls, OSA subjects had an elevated controller gain (control 0.24 ± 0.04 vs. OSA 0.36 ± 0.15; $p<0.05$) and a higher loop gain (control 0.07 ± 0.04 vs. OSA 0.14 ± 0.10; $p<0.05$). No significant differences between plant gains were observed (control 0.28 ± 0.12 vs. OSA 0.35 ± 0.13).

3.2.6 Discussion

The major contribution of the current study is the introduction of a chemoreflex system identification technique capable of adapting to changes in controller gain, plant gain, and loop gain. The proposed framework is non-invasive (i.e., based on spontaneous breathing). Through the use of a signal quality index, the deleterious effects of recording artifacts on the parameter estimation procedure are mitigated. The estimation technique tracked changes in chemoreflex parameters during physiological recordings of human subjects during sleep, where the controller gain was mechanically increased via PAV.

Breath-to-breath variability in ventilation has been studied extensively [63, 84], with the *periodic* component of variability generally attributed to the chemoreflex feedback mechanism, and the *nonperiodic* component ascribed to neural, vagal, behavioral, and other factors external to the chemoreflex feedback loop. Several authors have employed system identification techniques based

on autoregressive modeling that make use of the breath-to-breath dependence of ventilation on arterial blood gas tensions (via chemoreceptors) and dependence of O2/CO2 on ventilation (due to gas exchange) [33, 53, 71]. Other researchers used the autoregressive modeling in the presence of exogenous stimulations (e.g., pseudorandom binary sequences of inhaled CO_2) [32, 64, 85]. However, recordings of respiratory variables are subject to measurement artifacts, and the presence of hypopnea and apnea can result in degradation or loss of end-tidal CO_2 measurements. Moreover, externally administered gases can interfere with the subject's sleep pattern, thus complicating the theoretically appealing features of pseudorandom binary stimulation technique [86]. The proposed technique makes it possible to track the state of the control system under different experimental condition, or as a result of changes due to sleep state or body position.

OSA and Ventilatory Stability

Our findings suggest that OSA subjects have elevated controller gain compared to healthy controls, but no difference in plant gain, with an overall less stable control system, i.e., higher loop gain in OSA. Of note, Loewen *et al.* [87] and Salloum et al. [88] suggested that in untreated OSA patients a high LG is predominantly driven by an elevated controller gain. However, our observation of the differences between the two groups is confounded by the fact that our subjects were not BMI- and age-matched. Moreover, the OSA subjects were placed on a higher level of CPAP; however, CPAP is known to reduce plant gain by increasing lung volume, and therefore its overall effect would be a reduction in loop gain [70, 89] . Thus, we cannot conclude with certainty that OSA per se had higher loop gain in our sample, but our purpose was to assess our newer technique rather than to confirm the known elevation of loop gain seen in OSA.

Methodological Considerations

Although we selected periods of stable NREM sleep without evidence of clinical arousals, we cannot discount the influence of more subtle EEG changes. Future work should either model the influence of arousals on ventilation or include signal quality indices that minimize the contribution of large arousal-related ventilation on the estimation procedure.

Conclusion

Our results demonstrate that 1) the proposed autoregressive model is able to detect the expected increase in controller gain and overall LG with PAV and 2) that OSA subjects may have an elevated

controller and thus LG when compared to healthy controls. Such work strengthens the importance of LG in the pathogenesis of OSA and highlights that our model is an easy tool that could be used clinically for the identification of those who may benefit from interventions that lower LG (e.g., acetazolamide).

3.3 Appendix

3.3.1 Model Order Selection and Data Segmentation

Choice of Model Order

In autoregressive modeling, the optimal model order is typically estimated from the data through evaluation of Akaike's information criterion (AIC) [90]. The AIC selects a model that optimizes the trade-off between a parsimonious model (lowest model order) and one that provides the best predictive power (minimizing prediction error). The use of AIC is based on the assumption that no prior information concerning the model structure is available. Note that our proposed model is *physiologically structured*, based upon our current understanding of the physiology of the respiratory control system. Dependence of the key variables on their own history captures the notion of memory in the system. In the case of ventilation this memory is neural [91], while in the case of blood gases this memory may reflect the buffering effect of the functional residual capacity and lung tissue [65].

In this work, the orders of the autoregressive terms associated with the dynamics of the plant, the terms $a_{PCO2,\dot{V}}$, $a_{PCO2,PCO2}$, $a_{PO2,\dot{V}}$, $a_{PO2,PO2}$ in Eq. (3.1, were set equal to 1 in accordance with earlier studies [32, 65, 92]. Moreover, the terms $a_{PCO2,PO2}(p)$ and $a_{PO2,PCO2}(p)$ were set to zero for all p, since we assumed a negligible interdependence of $PCO2$ and $PO2$. Such an assumption vastly simplifies the model analysis. This is justified because our model deals with small 1^{st}-order fluctuations about nominal resting values; the interaction term is 2^{nd} order, and can therefore have only a minor effect on the measured ventilatory control variables. Although these choices of model parameters for the plant resulted in 9 out of 15 nonwhite residuals (that is, model residuals with statistically significant correlation with respect to the variable being estimated for lags 1 or 2), the identified plant dynamics is physiologically sound. Increasing the model order by up to two additional lags eliminated the issue of whiteness of residuals, but made the results harder to interpret, so those results are not discussed in this work.

To achieve a minimal model capturing the chemoreflex dynamics, the model order associated

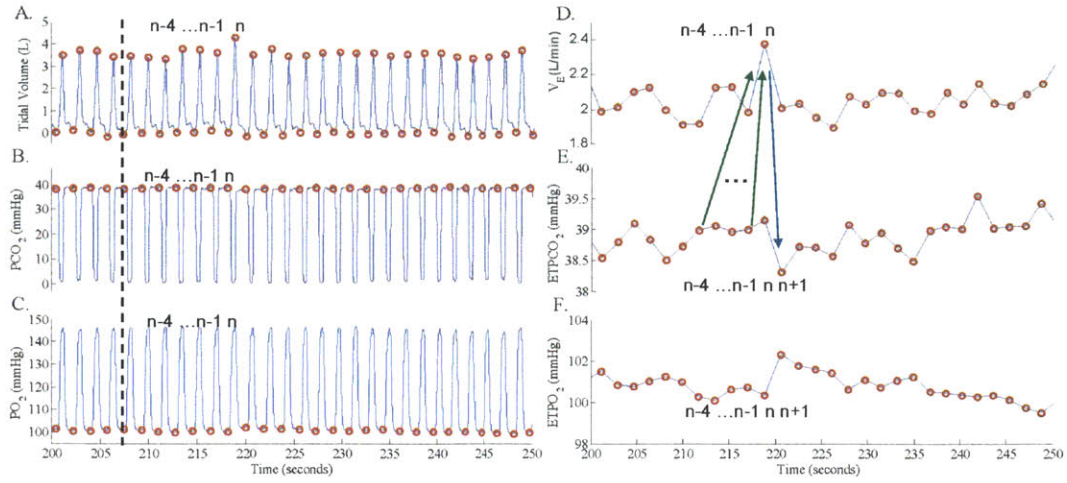

Figure 3-14: Polysomnographic waveforms.
An example of tidal volume, *PCO*2 and *PO*2 waveforms with marked fiducial points (o). The indices n, $n-1$, $n-2, \cdots$ reflect the breath-to-breath time-series alignment convention adopted for the purpose of this study. The double-ended arrow below the *PCO*2 waveform indicates the expiratory region for the associated breath. The mean over the last quarter segment of this region (indicated by a + sign) was taken as the end-tidal *PCO*2 value for that breath. Panel D graphically shows which values of *PCO*2 influence the current value of ventilation (green arrows), and similarly which *PCO*2 measurement is influenced by the current value of ventilation (blue arrow).

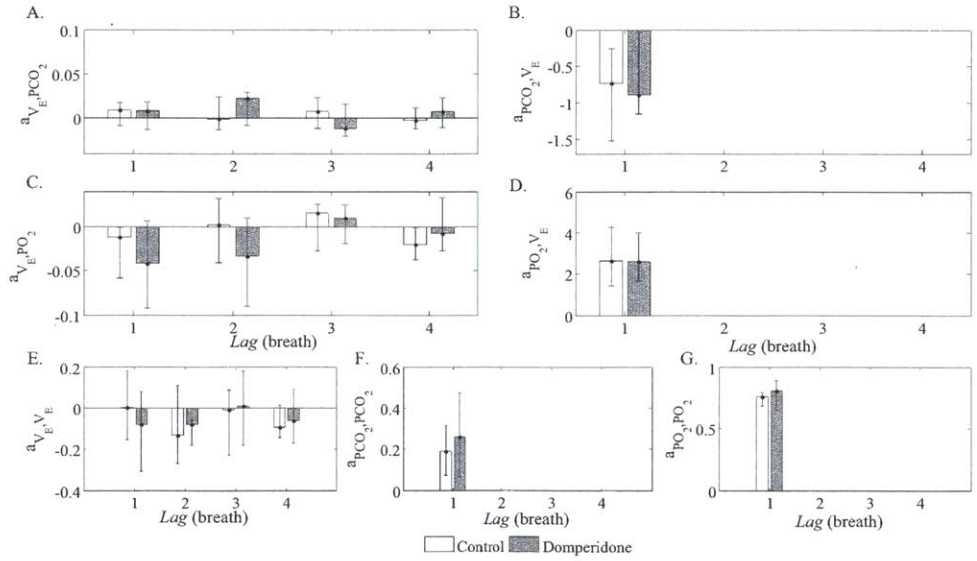

Figure 3-15: Group averages and standard deviations of the identified model parameters for the control (white) and domperidone (grey) conditions.
In each case, the horizontal axis represents the lag (the k index in Eq. 3.1) and the vertical axis is the amplitude of the regression between the associated signals. For example, the lag-2 coefficients $a_{\dot{V},PCO2}[2]$ (panel A) represents the influence of *PCO*2 three breaths earlier (2 lags and an extra lag due to the particular waveform alignment convention adopted in this work, see Fig. 3-14) on $\dot{V}$ at the current breath. This lag represents the convective delay for pulmonary capillary blood to reach the carotid chemoreceptors.

with the autoregressive terms $a_{\dot{V},\dot{V}}$, $a_{\dot{V},PCO2}$, and $a_{\dot{V},PO2}$ in Eq. (3.1) was increased sequentially from 2 to 10 breaths, and the lowest AIC order which resulted in white residuals was selected [34]. The obtained model orders varied between 2 and 5 time lags across subjects. From a physiological point of view, it takes at least 2-3 breaths for a change in blood gases at the alveolar level to reach the peripheral chemoreceptors. Therefore, to have a model that is physiologically sound and to facilitate intersubject comparison and group averaging, we set the model order equal to 4 for all the subjects and across different protocols. The model order utilized here is well within the bounds of the orders reported in univariate analysis of the chemoreflexes in Khoo and Marmarelis [53], also by Ghazanshahi and Khoo [32], who used autoregressive and moving average of order 1 (ARMA1), and 1-3 pure lags representing the convective delay for pulmonary capillary blood to reach the carotid chemoreceptors (thus considering signal history up to 4 lags); and the order reported by Clement and Robbins [93], who estimated a median lag of 2 breaths to carotid body chemoreceptors. Fig. 3-15 shows the averages and standard deviations of the model parameter values between each pair of variables.

3.3.2 Window Size Selection

The influence of window size on model fitting was studied by varying the window size (75, 100, 125, 150 breaths), and a window size of 125 breaths was selected based on the smallest AIC and loss function. The loss function is defined as the determinant of the residual covariance matrix Σ. The AIC is related to the loss function, model order (P), and window size (N) via the relationship $AIC = \log(\det(\Sigma)) + 2P/N$. In general, it is recommended to have several times as many data points as the number of model parameters (here we have 16 parameters to fit, so 125/16$\sim$ 8). On the other hand, the window size has to be kept small enough to ensure appropriate stationarity of the joint time-series [94]. The choice of 125 breaths as the window size resulted in 1-5 non-overlapping windows of spontaneous breathing data across different subjects. For each subject the window with the largest power in the medium frequency band of $\dot{V}$ (assessed via periodogram-based power spectral estimation) was selected for model fitting purposes. Our rationale behind this choice of data segment selection was to enhance model identification by choosing a segment of data with sufficient variability (see Mthodological Considerations under Discussion) and to characterize the system at its most variable (or unstable) observed state during spontaneous breathing.

Chapter 4

Discovery of Shared Dynamics in Multivariate Cohort Time Series

4.1 Introduction

This chapter examines the problem of segmenting physiological time-series within a patient cohort into stereotypical (or phenotypical) dynamics, with the objective of extracting informative features for predicting patient outcome variables of interest. We present the class of switching linear dynamical systems (SLDS), as an alternative approach to modeling nonlinear and nonstationary time-series, when the system can be assumed to be piecewise linear in time. We provide a high-level overview of the the expectation-maximization (EM) algorithm for parameter estimation of the SLDS models, and point out some of the shortcomings of EM, depending on the objectives of time-series modeling. An alternative "outcome-discriminative" learning algorithm in presented in the next chapter for the case when SLDS models are used for extraction of useful features from the data with the aim of making predictions. A preliminary version of this chapter was presented at the 34th annual international conference of the IEEE engineering in medicine and biology society (EMBC'12) [95].

4.1.1 Background

Multivariate autoregressive (MVAR) models of physiological time-series have been used by several authors to infer properties (e.g., transfer functions) of the underlying dynamical systems [25,30,33]. For instance, several authors have used the time-series of heart rate (HR) and blood pressure (BP) to obtain estimates of baroreflex gain [26, 96, 97], or more generally to characterize the feedback

control system regulating the cardiovascular variables [24]. The linear techniques commonly used (often based on variants of autoregressive modeling) have the advantage of revealing the individual relationships among the observed variables (e.g., the baroreflex gain describes the relationship between HR and BP, excluding the possible influence of respiration). On the other hand, nonlinear indices of complexity are capable of capturing a richer set of dynamical behavior [28,98], but often lack physiological interpretability in terms of specific underlying mechanisms.

In this work, we assume that although the underlying dynamical system may be nonlinear and the stochastic noise components can be non-Gaussian, the dynamics can be well approximated by a mixture of linear dynamical systems. We will refer to the components of such a mixture as "modes". Specifically, we present a technique based on a switching linear dynamical system [42] that is: (i) sufficiently simple to allow for a physiologically-interpretable model of the interaction between HR and BP, (ii) sufficiently complex to provide a realistic representation of the underlying physiology, and (iii) provides a framework for defining a measure of similarity among multiple or cohort multivariate physiological time-series based on their underlying shared dynamics.

The mixture modeling approach provides a framework for automatic segmentation of time-series into regions with similar dynamics (i.e., time-dependent rules describing how the future state of the system evolves from the current state). Furthermore, we assume that similar subjects respond similarly to perturbations (such as tilting), and therefore share dynamical modes. The latter assumption allows us to define a notion of similarity across segments from multiple time-series. This approach provides a potential improvement over time-series similarity measures that are based on symbolic representations [99] or simple trend detection [100]. These measures often ignore the joint temporal information that is embedded in the dynamics of interaction among physiological variables.

A central aim of the current work is to develop a framework for automated discovery of evolving dynamics in multivariate physiological time-series from large patient cohorts, such as the Multi-parameter Intelligent Monitoring for Intensive Care II (MIMIC II) database of over 30000 patients [29]. Intensive care units (ICUs) are some of the most important components of the health care system. With the ubiquity of inexpensive high-capacity recording and storage devices, it is becoming possible to continuously record and archive patient vital signs, such as heart rate (HR) and blood pressure (PB) [29]. Despite this continuous feed of data, the commonly used acuity scores, such as APACHE and SAPS, are based on snap-shot observations of the patient [101–103]. However, physiologic systems generate complex dynamics in their output signals that reflect the

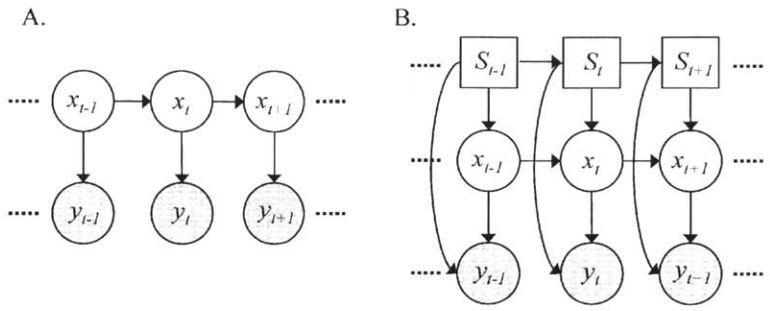

Figure 4-1: Graphical model representation.
Panels A and B show the graphical model representations of the state-space and the switching state-space models, respectively. The round nodes are continuous and Gaussian, and the square nodes are discrete. Shaded nodes are observed and the rest are hidden. Arrows denote the conditional dependencies among the random variables. A time-unrolled detailed representation of the switching state-space model of panel B is shown in Fig. 5-2.

changing state of the underlying control systems [27, 28, 33]. For instance, time-series of BP can exhibit oscillations on the order of seconds (e.g., due to the variations in sympathovagal balance), to minutes (e.g., as a consequence of fever, blood loss), to hours (e.g., due to sleep-wake cycle and circadian effects). Discovering and understanding these dynamical behaviors are of both fundamental and clinical importance [104].

4.2 Modeling Switching Dynamics in Cohort Time Series

The time-varying MVAR identification approach discussed in Chapters 2 and 3 (Part II) requires estimating as many model coefficients as the number of observations, although the adjacent estimates are generally highly correlated (depending on settings of the D matrix). This over-parametrization may result in a large estimation variance. An alternative approach is to assume a small set of linear models, and either switch between them or to use a weighted combination of them (i.e., a mixture model) to explain the evolving dynamics of the system. The switching linear dynamical systems (SLDS) framework [42] utilizes a bank of J different linear dynamical models, along with a set of switching variables with a Markov transition dynamic that model the probability of belonging to any one of the J modes (see Fig. 4-1). In particular, when the objective of time-series modeling is feature extraction with the aim of predicting outcome variables of interest, the SLDS approach requires fewer parameters (and therefore, allows a more compact representation of the dynamics).

4.2.1 Switching Kalman Filter

Let us define a mode (or dynamic) as a set of model parameters $\Theta = \{A,\ C,\ Q,\ R\}$ (also known as a dynamical mode). We further assume a set of (switching) variables S_t that specify the respective probabilities of each of the J modes being active at any given time t. The switching Kalman filter (SKF) algorithm [42] assumes S_t follows a Markov chain with $(J \times J)$ transition matrix Z and $(1 \times J)$ initial distribution vector π.

4.2.2 EM for Parameter Learning in Switching Dynamical Systems

In practice, we neither know the set of switching variables (i.e., segmentation of the time-series) nor the modes. The EM algorithm for finding the maximum-likelihood set of model parameters has been briefly discussed in Chapter 2, Section 2.3.2; for a more comprehensive treatment of EM in the context of SLDS models see Murphy (1998) [42]. Here we highlight the main intuition behind the EM algorithm, as well as our modifications (1) to impose physiological constraints on the dynamics for each mode, and (2) to learn shared modes across all the time-series within a patient cohort.

Consider a set of N patients with time-series $y_t^{(n)}$ of length T_n. We choose a random initialization of $\Theta^{(j)}$ as well as $x_{0|0}^{(j)}$, $V_{0|0}^{(j)}$), Z, π. The E-step of the EM algorithm involves constructing an approximating distribution over the hidden states $x_t^{(n)}$ and the corresponding switching variables $S_t^{(n)}$, using a modified Kalman smoother [42]. We run the E-step separately on each of the N time-series. Next, given the observations $y_t^{(n)}$, and the approximate distributions over $x_t^{(n)}$ and $S_t^{(n)}$, the M-step maximizes the expected complete data log-likelihood by adjusting the model parameters across all the modes (via constrained least-squares optimization to maintain the structure of the A matrices, as discussed in Chapter 2, Section 2.5.3), as well as the Markov chain parameters Z and π. This learning of the model parameters is done by pooling together all the observed and inferred hidden variables across *all* the subjects. Iteration through several steps of the EM algorithm will result in learning a set of J shared modes and a global transition matrix Z for all the patients. In addition, to minimize the influence of initialization on the final learned parameters and to reduce the chance of trapping in local minima, one may implement a deterministic annealing step within the EM algorithm [42]. Hereafter, for simplicity, we assume that the Q and R matrices are diagonal.

4.2.3 Switching Dynamical Systems for Feature Extraction and Prediction

Let us define a "mode proportion" (denoted by η_j) as the proportion of time a time-series (or a patient) spends within the j-th mode. Given the maximum expected log-likelihood estimates of the switching variables S_t from the EM algorithm (for each time-series/patient), we have

$$\eta_j = \frac{1}{T}\sum_{t=1}^{T} M_t^s(j) \tag{4.1}$$

where $M_t^s(j) = Prob(S_t = j|y_{1:T})$. Note that, $\sum_j \eta_j = 1$. For subsequent classification and prediction purposes, we replace each time-series with its corresponding mode proportion (a compact $1 \times J$ feature-vector), and use a logistic regression classifier to make predictions about the outcome variables of interest.

4.3 Datasets

4.3.1 Cardiovascular Simulation

We simulated time-series of cardiovascular control system variables based on a delay recruitment model of the cardiovascular control system, as described in Fowler and McGuinness [105] and McSharry et al. [106]. The model included a coupled system of nonlinear delayed differential equations, controlling HR and BP, with respiration as an exogenous input. We simulated 10 different multivariate time-series of HR and mean arterial BP (MAP), each including three different dynamics that become dominant in a random order and last for a variable length of time. The three dynamics (color-coded as red, blue, and black, respectively, in Fig. 4-2) approximate aging-related autonomic changes: a progressive reduction in parasympathetic gain (from 0.40 to 0.13 to 0.07 in normalized units, see [105]) and an increase in sympathetic delay (from 3 to 5 seconds).

4.3.2 Tilt-Table Experiment

Time-series of HR and MAP were acquired from 10 healthy subjects undergoing a tilt-table experiment. The details of the protocol are described in Heldt et al. [16]. Briefly, subjects were placed in the supine position and secured to the table. Tilting was performed from the horizontal position to the vertical position and back to supine. Examples of the resulting time-series are shown in Fig. 4-3.

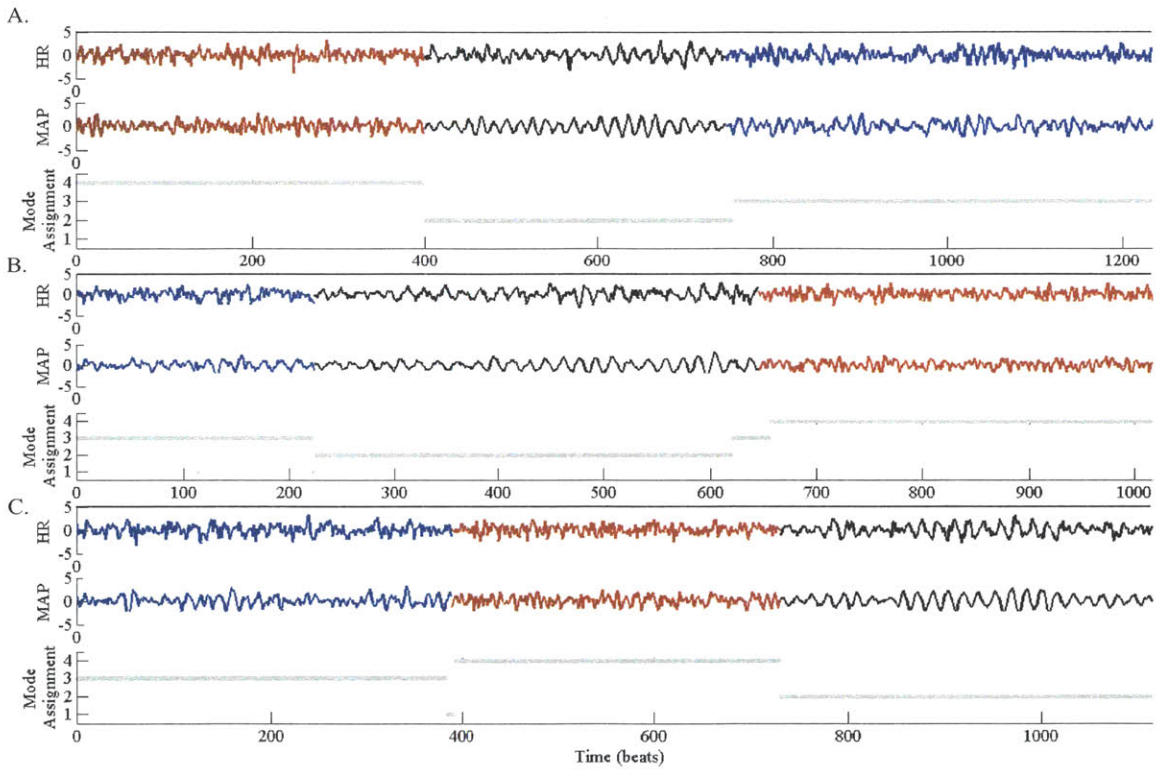

Figure 4-2: Simulation study of the cardiovascular system.
Three examples (out of the 10 simulated time-series) of HR and MAP (after filtering) are shown in panels A, B and C. In each case, the actual dynamics are color coded. The horizontal gray lines show the inferred most likely segmentations based on the switching Kalman filter (SKF). Note that, the SKF consistently assigned modes 2, 3 and 4 to the dynamics color-coded as black, blue and red, across all the simulated time-series.

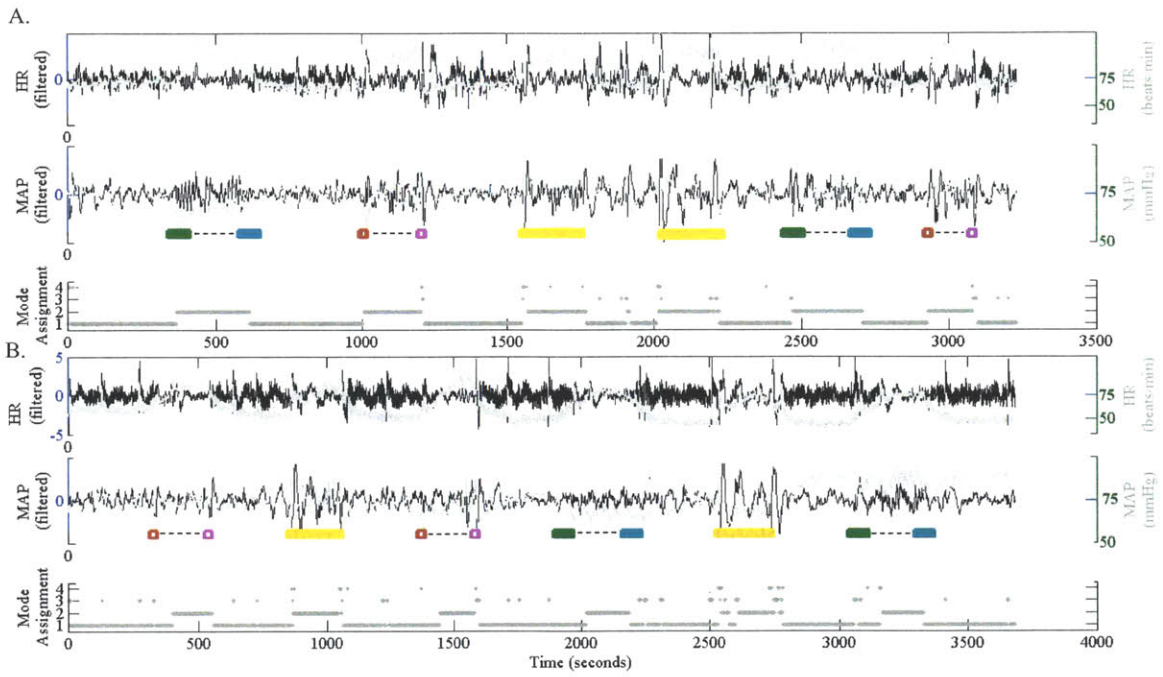

Figure 4-3: Tilt-table study.

Two examples out of the 10 recordings of HR and MAP from the tilt-table experiment are shown in panels A, B. Within each panel, from top to bottom, HR and MAP (actual values in gray and filtered values in black), and SKF-based segmentation are shown. In each case, the actual true dynamics are color coded (green to cyan: slow tilt up and down to supine; red to pink: rapid tilt up and down to supine; yellow: standing up and back to supine). Note that the SKF consistently assigns the modes 1 and 2 to the supine and non-supine states, respectively. The other two modes seem to capture the high-frequency noise components of the time-series.

4.3.3 MIMIC Database of Intensive Care Unit Patients

The MIMIC II waveform database (version 2) [29] includes approximately 4,000 sets of high resolution physiological waveforms with associated minute-by-minute vital sign trends. This study used only the adult patients from the MIMIC II waveform database with clinical information, and with at least 8 hours of continuous minute-by-minute invasive BP trends during the first 24 hours of their ICU stays. Patients with more than 15% of missing or invalid (i.e., outside physiologically plausible bounds of 20 to 200 mmHg for mean pressures) BP samples were excluded, yielding a total of 479 patients. The remaining missing or invalid values were replaced with random Gaussian noise (zero mean, and variance one), which should correspond to a single mode in the final learned model . Of all the patients within the selected cohort, 16% died before hospital discharge. The dataset contained approximately 9,700 hours of minute-by-minute systolic BP measurements (20.2 hours per patient on average). In order to compare with the SAPS-I score, we restricted our regression analysis to 452 patients whose records included SAPS-I score during the first 24 hours of their ICU stays. SAPS-I uses a weighted combination of several medical attributes to assess patient mortality. These attributes include the most extreme values of systolic blood pressure, heart rate, temperature, blood urea nitrogen, white blood count, serum potassium, serum sodium, serum bicarbonate, serum bilirubin, as well as age, presence/absence of leukocytosis, presence/absence of acidosis and presence/absence of oliguria.

4.3.4 Data Pre-processing

Since we are interested, for both the simulated data and the tilt-table experiment, in the dynamics of interaction between HR and MAP in the frequency range pertinent to sympathetic and parasympathetic regulation [107], the time-series of HR and MAP were high-pass filtered in both studies to remove the steady-state baseline and any oscillation in the time-series with a period slower than 100 beats. This filtering was done using a 7th order Butterworth digital filter with cutoff frequency of 0.01 cycles/beat. In the case of the MIMIC II database, the minute-by-minute time-series of blood pressure were filtered to remove linear trends.

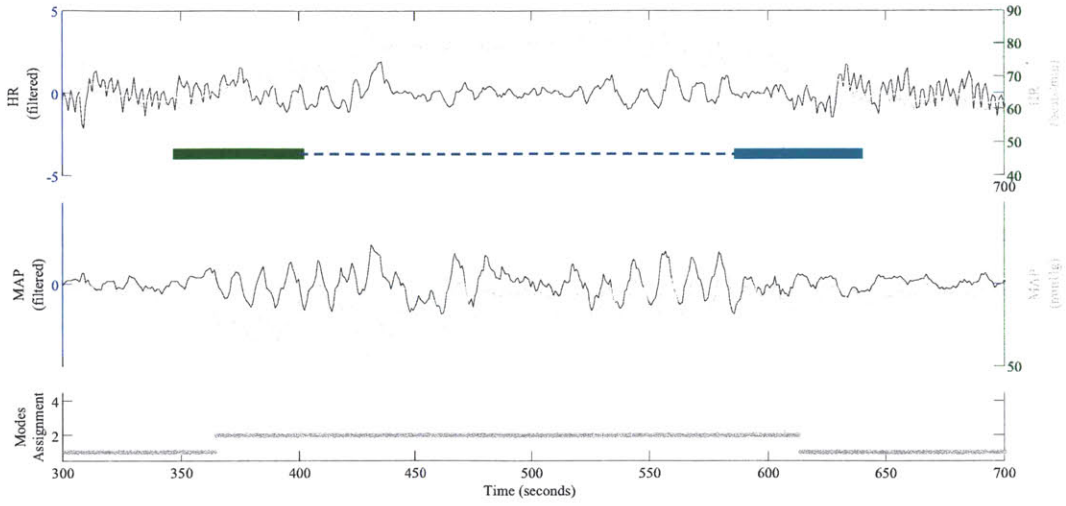

(a) Slow Tilting

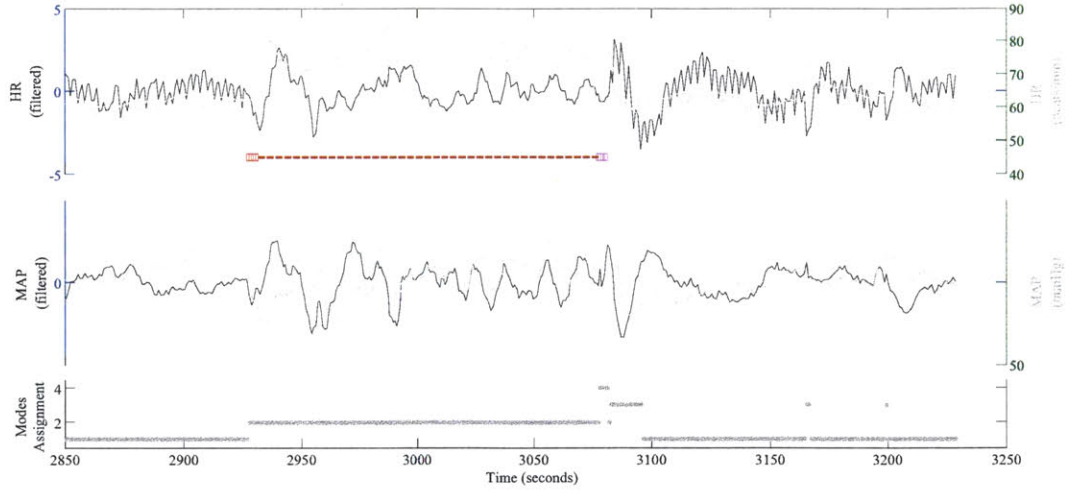

(b) Fast Tilting

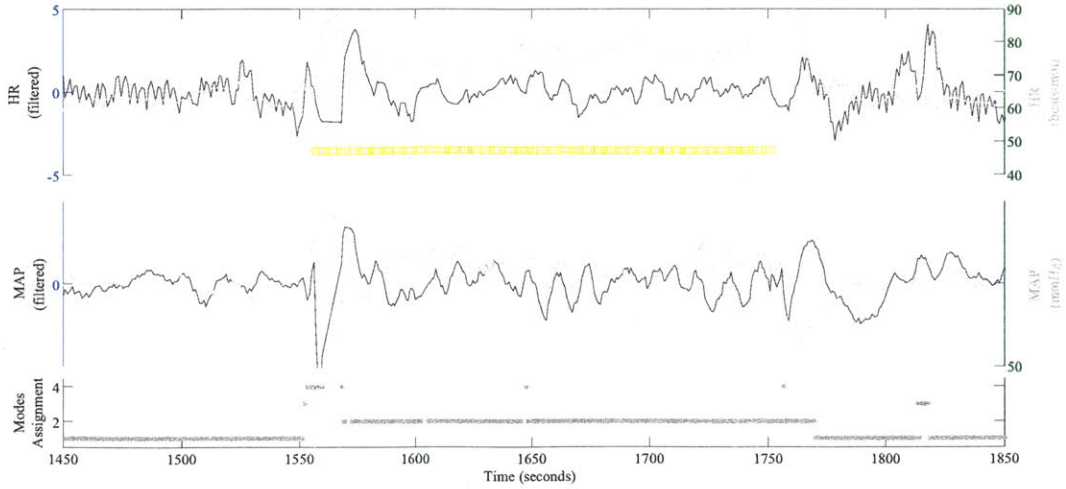

(c) Standing up

Figure 4-4: Example time-series from the tilt-table experiment.

4.4 Results

4.4.1 A Simulated Illustrative Example

Fig. 4-2 A-C shows three examples of simulated time-series with the inferred SKF-based segmentations. In all 10 simulated cases the SKF algorithm was able to find the correct segmentation of each time-series, as well as the sharing of the dynamics across multiple time-series.

4.4.2 Tilt-Table Experiment

Figure 4-3 shows time-series of HR and MAP from two subjects, as they undergo the different tilting protocols. 4-4 shows sample time-series corresponding to the slow (panel A), fast (panel B), and standing (panel C) dynamics of HR and MAP.

We constructed a binary classification problem of determining supine vs. non-supine segments (see Fig. 4-3). There were a total of seven supine positions and six non-supine positions per subject (total of 13×10=130 binary outcomes). We used the average probability of belonging to each of the four SKF states (J=4) as our feature-vector (i.e., average value of $P(S_t)$ within each position), resulting in a 130×4 input matrix. Application of logistic regression with 10 fold cross-validation yielded an AUC of 0.97 ± 0.01, indicating an excellent discriminatory power.

4.4.3 Time-series Dynamics and Hospital Mortality

In the case of the MIMIC II dataset our aim was to use the mode proportions associated with the minute-by-minute mean blood pressure time-series dynamics to predict probabilities of in-hospital patient mortality. The performance of the proposed classification task is reported using 10 fold cross-validated AUC (median [interquartiles]) as follows. For each fold the cohort was divided into a training dataset (90% of the subjects) and a testing dataset (10% of the subjects) with a similar proportion of survived and expired patients. The results presented here are based on using 20 modes, with the subsequent addition of two more features: SAPS-I score and ICU care unit.

In-Hospital Mortality Prediction

Table 4.1 summarizes the performance of the SAPS-I acuity score, the mode proportions, as well as their combined performance with the addition of the ICU care unit. The median hospital

	SAPS-I	MP	MP&SAPS-I	CU&SAPS-I	MP&CU	MP&SAPS-I&CU
AUC	0.66 [0.62 0.70]	0.67 [0.64 0.70]	0.73 [0.69 0.75]	0.69 [0.56 0.73]	0.72 [0.68 0.74]	0.77 [0.74 0.78]

Table 4.1: In-hospital mortality prediction (10 fold cross-validated).
Blood pressure dynamics (MP), SAPS-I score, and care unit (CU) were used as feature vectors. These results indicate that SAPS-I score and MP have independent predictive value.

	N (Mort %)	SAPS-I	MP	MP & SAPS-I
CCU	97 (16.5%)	0.68	0.72	0.76
CSRU	208 (4.3%)	0.58	0.75	0.75
MICU	119 (30.3%)	0.59	0.58	0.61
SICU	55 (21.8%)	0.86	0.57	0.75

Table 4.2: In-hospital mortality prediction, broken down by care unit.
CCU (coronary care unit), CSRU (cardiac surgery recovery unit), MICU (medical intensive care unit), SICU (surgical intensive care unit).

mortality prediction using the mode proportions was similar to the SAPS-I score (0.67 vs. 0.66). Moreover, combining features from the mode proportions and SAPS-I score resulted in a significant improvement in performance (0.73). Notably, adding the care unit as a feature improved the performance of both SAPS-I and mode proportions, suggesting that the two features may have a more robust performance in certain care units.

To further explore these observations, we pooled together all the subjects in the testing sets of all 10 folds, and calculated the care unit-specific AUCs. The results presented in Table 4.2 demonstrate a better performance of the mode proportions in two of the care units (CCU, CSRU), while the SAPS score performs exceptionally well in the SICU. (A possible explanation for the latter observation is that in preparation for surgery accurate lab and other measurements are obtained, which could enhance the performance of the SAPS-I score. However, further studies on the entire MIMIC cohort are needed to confirm this finding.) In summary, these results suggest that within three out of the four care units the addition of mode information may improve the ability to assess patient risk and mortality. Moreover, the superior performance of the mode proportions within the CCU and CSRU may be due to fact that patients within these units suffer primarily from cardiac related illnesses, and therefore their blood pressure dynamics is informative of the state of their health.

	SAPS-I	MP	MP&SAPS-I	CU&SAPS-I	MP&CU	MP&SAPS-I&CU
AUC	0.59	0.70	0.71	0.68	0.72	0.73
	[0.56 0.66]	[0.63 0.73]	[0.69 0.77]	[0.53 0.73]	[0.69 0.73]	[0.70 0.78]

Table 4.3: Thirty-day mortality prediction (10 fold cross-validated).
Blood pressure dynamics (MP), SAPS-I score, and care unit (CU) were used as feature vectors. These results indicate that SAPS-I score and MP have independent predictive value.

	N (Mort %)	SAPS-I	MP	MP & SAPS-I
CCU	97 (21.6 %)	0.62	0.81	0.84
CSRU	208 (4.3 %)	0.57	0.77	0.76
MICU	119 (37.8 %)	0.56	0.61	0.63
SICU	55 (25.45 %)	0.86	0.41	0.58

Table 4.4: Thirty days mortality prediction, broken down by care unit.

Thirty-Day Mortality Prediction

To explore the utility of blood pressure dynamics for long-term prognosis of survival and mortality, we repeated the study of the previous section using a 30-day mortality window. Out-of-hospital mortality was assessed by looking up social security records. Tables 4.3 and 4.4 summarize the results of 30-day mortality prediction. Notably, the mode proportions alone achieve 81% accuracy within the CCU. These result suggest that the dynamics of blood pressure within the cardiac intensive care unit are a significant predictor of a patient's overall well-being.

4.5 Discussion and Conclusion

We presented a technique for extracting shared physiological dynamics within a cohort. It was shown that discrete regime changes in the dynamics of HR and BP, either as a result of an altered underlying control system (e.g., increased sympathetic regulation with tilting) or due to external perturbations (e.g., response to drugs and interventions in ICU), can be captured in an automated fashion. We showed that the discovered dynamics can be used for prediction and tracking of a patient's propensity to survive the ICU stay. Interestingly, the BP time-series dynamics alone had a comparable performance to that of the SAPS-I score, although the latter included 14 lab measurements and demographic information. Moreover, our results indicate that the dynamics of the BP time-series, when combined with traditional patient acuity scores and information regarding the care units, provides a more accurate assessment of patient survival/mortality in the ICU. Notably,

we observed that within two cardiac care units (CCU and CSRU) the blood pressure dynamics alone were a significant predictor of outcomes for both in-hospital and 30-day out-of-hospital mortality. There results are complementary to the observations of Celi et al. [108], who reported that customization of mortality prediction techniques to acute kidney injury patients may result in significant improvement in prognostic performance.

Since an AR model of HR and BP regulation was used, the proposed framework allows us to extract useful indices of baroreflex activity from the learned AR coefficients [24, 33]. Therefore, the discovered dynamical patterns are physiologically interpretable. Future work should involve exploring the full potential of the SLDS framework to model nonlinear dynamics and non-Gaussian physiological and measurement-related noise. The technique is most useful in modeling systems with rapid transition among physiological states (e.g., study of changes in autonomic regulation with sleep stages).

Other approaches to inference and learning within the switching dynamical system frame work include the nonparametric Bayesian approach and other sampling based methods [109, 110]. The current framework has two main advantages. First, the inference (i.e., latent state estimation) step of the algorithm can be parallelized (i.e., performed separately for each patient) and processed on a cluster computer, thus allowing the technique to scale up to thousands of patient time-series. (Note, the nonparametric BP-AR-HMM technique of Fox et al. (2009) [109] is not amenable to parallel processing due to the way inference is performed in the conventional implementation of the beta-process). Secondly, the switching Kalman filter technique for inference within the switching linear dynamical system framework has the added advantage of having a closed-form (approximate) analytical solution. This in turn allows for designing an efficient outcome-discriminative (or supervised) learning algorithm, where the learning of dynamics can be done with the goal of making predictions about outcome variables of interest, as will be discussed in the next chapter.

Chapter 5

Learning Outcome-Discriminative Dynamics in Cohort Time Series

5.1 Introduction

Extending the switching linear dynamical systems (SLDS) framework discussed in the previous chapter, we propose a new class of learning algorithms for "outcome-discriminative" discovery of dynamics within cohort time-series. An example will help to motivate the need for an alternative approach for learning SLDS differently than via the traditional maximum likelihood-based techniques (such as the EM algorithm). Consider the simulated time-series shown in Fig. 5-1, where four bivariate time-series, each having different proportions of four different modes, are shown. Let us assume that we are given a collection of such time-series, and asked to classify the time-series into four categories. Common to all time-series is the presence of two randomly placed large amplitude artifacts (uniform random noise). A maximum likelihood-based learning algorithm may use one or more modes to explain such artifacts. However, since the artifacts are common to all four categories, and thus are not relevant to the outcomes, learning of these artifactual patterns will not improve our classification performance.

5.2 Outcome-Discriminative Learning

The objective of the development in this chapter is to design an outcome-discriminative learning algorithm that discovers time-series features relevant to outcome variables of interest. The approach we take is based on exploiting a relationship between probabilistic dynamic Bayesian networks

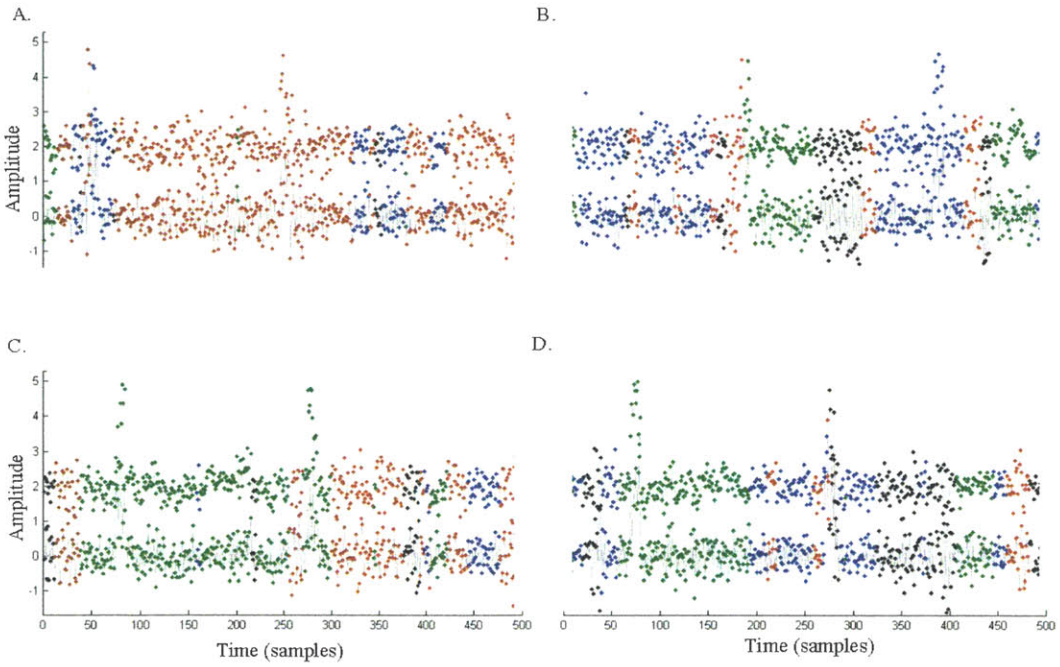

Figure 5-1: Example simulated time-series from four different categories.
Four example bivariate time-series corresponding to four latent categories A, B, C, and D (see Section 5.7.1 for simulation details) are shown in their respective panels. Each category includes different proportions of four different modes. The four modes are color-coded (red=1; blue=2; green=3; black=4). All four categories include two randomly placed large amplitude artifacts (uniform random noise in the interval of [0 15]) of 10 samples duration. (Note, we introduced an offset of 2 in one of the channels of each time-series to improve visibility.)

(such as the SKF) and deterministic neural networks, to perform efficient error gradient calculations for supervised learning of dynamics. Figure 5-2 presents a schematic diagram of a neural network representation of the switching state-space model presented in Fig. 4-1, augmented with a binary logistic classification layer. In general, the classification layer may take the form of a multinomial Logit or any multi-layer neural network, and the associated cost functions may include additional regularization terms. In addition, one may construct more general feature vectors, based on the marginal distributions of the latent variables.

Given the representation of Fig. 5-2, one may learn the marginal distributions (corresponding to the nodes) and the model parameters using the standard backpropagation technique for neural networks, with the goal of maximizing the probability of outcomes (or minimizing the prediction error). The objective of error backpropagation is to calculate the sensitivity of a given error function with respect to small variations in the model parameters. To accomplish this, the backpropagation technique uses the *chain rule* (from calculus) to efficiently calculate the error gradients with respect

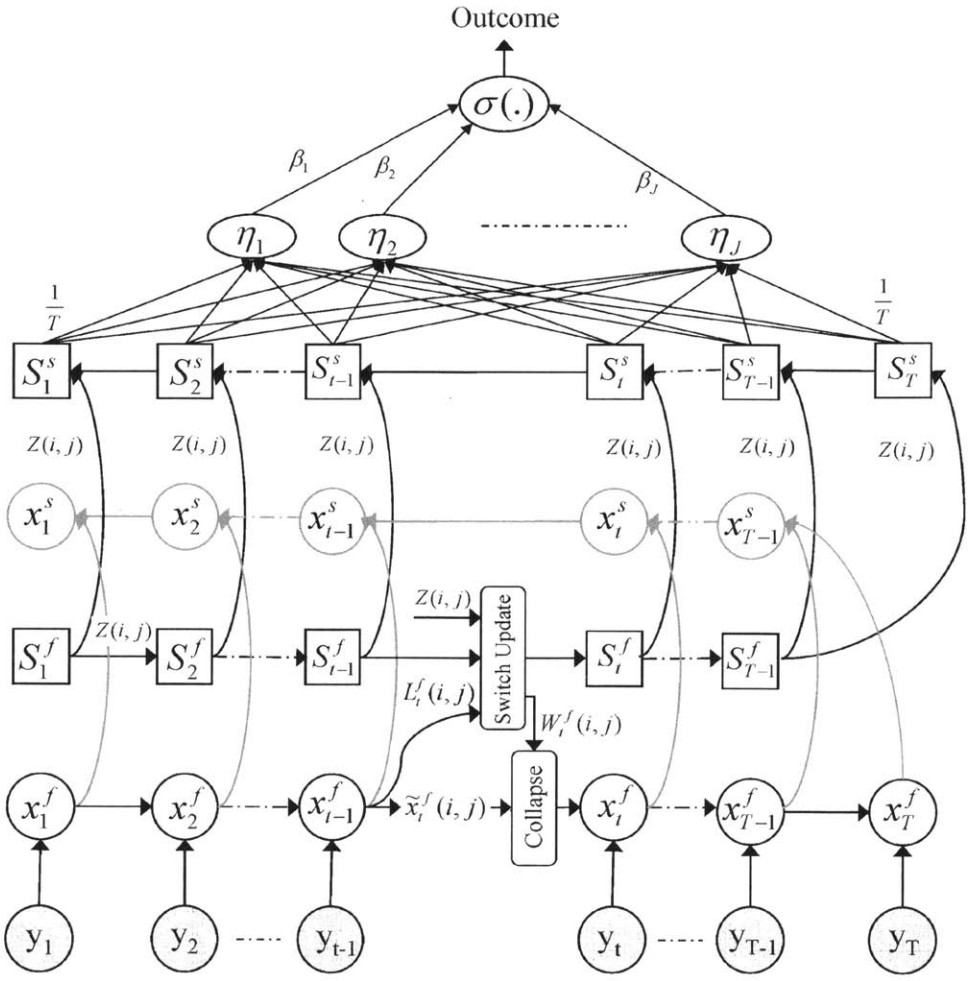

Figure 5-2: Information flow in a switching linear dynamic system with an added logistic regression layer.

The x nodes contain M continuous Gaussian random variables, representing the state of M different modes. The $\tilde{x}$ nodes contain M^2 intermediate continuous Gaussian random variables, which are subsequently collapsed to M Gaussian random variable via *moment matching* [42]. Note that according to the SKF approximate inference algorithm, the smoothed switching variables (and therefore the mode proportions η_j) are independent of the smoothed state estimates (grey circled nodes x^s).

to all the marginal distributions and the model parameters. For the continuous Gaussian nodes we will need to calculate the gradients with respect to the means and covariances of each node.

Here we provide the *analytical* expressions for the involved derivatives, which allows a two-pass algorithm for calculating the exact gradients. The forward pass is equivalent to the inference step of the SKF, where we also estimate the partial derivatives of each node output with respect to its inputs. The backward pass uses the chain rule to recursively estimate the gradient of the error function with respect to the marginal distributions at each node, and the model parameters. Note that numerical estimation of gradients involves tweaking each of the model parameters and observing the change in the error function. For a model with hundreds to thousands of parameters, this amounts to running the inference step of the SKF many times. This approach is impractical for cohorts including hundreds to thousands of time-series, with hundreds to thousands of samples per time-series.

In what follows, we start with the regression layer shown in Fig. 5-2, and recursively estimate the error gradients all the way down to the filtering layer, and ultimately calculate the error gradient with respect to the parameters of the dynamical modes.

5.3 Derivatives of the Regression Layer

Logistic regression is commonly used for predicting outcomes of categorical variable. For instance, each of the N time-series within a cohort may be associated with an outcome (or label) denoted by $\{O_n^{true}\}_{n=1}^{N}$). In this work, we use the logistic and multinomial regression methods to map the mode proportions to the outcome variables of interest. Therefore, we provide the analytic gradients of the corresponding error functions with respect to the regression parameters and the mode proportions.

5.3.1 Binary Outcomes

Given a set of predictor variables (taken to be the mode proportions here), the binary logit function takes the following form:

$$\sigma(\bar{\eta}) = \frac{1}{1+exp(-\bar{\eta})} \tag{5.1}$$

with $\bar{\eta} = \beta_0 + \beta_1\eta_1 + \cdots + \beta_J\eta_J$. In Eq. (5.1), $\sigma(\bar{\eta})$ can be interpreted as the probability of a positive outcome, given the mode proportions (or more generally, model parameters). The Bernoulli

probability of a given outcome O^{true} ($\in \{0,1\}$) is then parameterized by $\sigma(\eta)$ we follows:

$$P(O^{true}|\sigma(\bar{\eta})) = \sigma(\bar{\eta})^{O^{true}}(1-\sigma(\bar{\eta})^{1-O^{true}} . \tag{5.2}$$

The objective of parameter fitting (finding the β_j coefficients) is to minimize the negative log-likelihood of the outcomes, given the predictor variables:

$$E = -logProb(O^{true};\ \bar{\eta}) = -(O^{true}log(\sigma(\bar{\eta})) + (1-O^{true})log(1-\sigma(\bar{\eta}))). \tag{5.3}$$

Note that the overall error is the sum over the individual errors, that is: $E_{all} = -\sum_{n=1}^{N} logProb(O_n^{true};\ \bar{\eta}^{(n)})$.

Error Gradient

The derivative of the error function in Eq. (5.3) with respect to $\bar{\eta}$ is given by

$$\frac{\partial E}{\partial \bar{\eta}} = (-\frac{O^{true}}{\sigma(\bar{\eta})} + \frac{1-O^{true}}{1-\sigma(\bar{\eta})})\frac{exp(-\bar{\eta})}{(1+exp(-\bar{\eta}))^2}. \tag{5.4}$$

Moreover, it follows from the Chain Rule that

$$\frac{\partial E}{\beta_j} = \frac{\partial E}{\partial \bar{\eta}}\eta_j , \tag{5.5}$$

$$\frac{\partial E}{\eta_j} = \frac{\partial E}{\partial \bar{\eta}}\beta_j . \tag{5.6}$$

5.3.2 Multinomial Outcomes

The multinomial probability of a given outcome O^{true} ($\in \{1,\cdots,K\}$), parameterized by μ_k, is given by

$$Prob(O^{true}|\mu_k) = \prod_{k=1}^{K} \mu_k^{O_k^{true}} = exp(\sum_{k=1}^{K} O_k^{true} log\mu_k) \tag{5.7}$$

where $\mu_k = \frac{exp(\theta_k)}{\sum_{k'} exp(\theta_{k'})}$, and $\theta_k = \beta_{k,0} + \sum_{j=1}^{J} \beta_{k,j}\eta_j$.

We take the error function to be the negative log-likelihood of the probability of an outcome given μ:

$$E = -logProb(O^{true}|\mu) = -\sum_{k=1}^{K} O_k^{true} log(\mu_k). \tag{5.8}$$

Note that the overall error is $E_{all} = -\sum_{n=1}^{N} logProb(O_n^{true}|\mu_k^{(n)})$.

Error Gradient

The error gradients with respect to the parameters $\beta_{k,j}$ of the multinomial regression and the mode proportions η_j are given by

$$\frac{\partial E}{\partial \beta_{k,j}} = \frac{\partial E}{\partial \theta_k} \eta_j \tag{5.9}$$

$$\frac{\partial E}{\partial \eta_j} = \sum_{k=1}^{K} \frac{\partial E}{\partial \theta_k} \beta_{k,j} , \tag{5.10}$$

where

$$\frac{\partial E}{\partial \theta_k} = -\sum_{k'=1}^{K} O_{k'}^{true} (\delta_{k,k'} - \mu_k). \tag{5.11}$$

5.4 Derivatives of the Switching Kalman Filter

Some Notation

Hereafter we will use parentheses to index individual modes, and we will use brackets to indicate the individual elements of a vector or a matrix. Thus, $A(i)[m,n]$ refers to the m-th row, n-th column element of the matrix of state dynamics for the i-th mode. We will use the symbol $\odot$ to denote the Frobenius inner product of two matrices (or vectors) defined as $A \odot B = \sum_i \sum_j A_{ij} B_{ij}$. Moreover, for a matrix B, indexed by i, j, the colon notation $B(i,:)$ denotes entries ranging over all j values. All the exponentiations involving matrices are element-wise. Finally, $\delta_{m,n}$ denotes a conformable matrix with the (m,n)-th element equal to one and zero elsewhere. Similarly, $\delta_{m,n}^{n,m}$ denotes a conformable matrix with ones at the (m,n)-th and (n,m)-th elements and zero elsewhere.

5.4.1 Filtering Step

For completeness we summarize the essential equations (pertinent to error backpropagation) of the filtering and smoothing steps of the SKF algorithm, as described in Murphy (1998) [42].

The filtering step of the SKF algorithm takes the following form [41,42]:

$$(\tilde{\mu}_{t|t}(i,j), \tilde{V}_{t|t}(i,j), L_t(i,j)) = KalmanFilter(\mu_{t-1|t-1}(i), V_{t-1|t-1}(i), \mathbf{y}_t(i); A(j), C(j), Q(j), R(j))$$

$$\begin{aligned}
\mu_{t|t-1} &= A\mu_{t-1|t-1} , && (5.12)\\
V_{t|t-1} &= AV_{t-1|t-1}A^{\mathsf{T}} + Q , && (5.13)\\
\mathbf{e}_t &= \mathbf{y}_t - C\mu_{t|t-1} , && (5.14)\\
\mathscr{S}_t &= CV_{t|t-1}C^{\mathsf{T}} + R , && (5.15)\\
G_t &= V_{t|t-1}C^{\mathsf{T}}\mathscr{S}_t^{-1} , && (5.16)\\
\tilde{\mu}_{t|t} &= \mu_{t|t-1} + G_t\mathbf{e}_t , && (5.17)\\
\tilde{V}_{t|t} &= (I - G_tC)V_{t|t-1} , && (5.18)\\
L_t &= N(\mathbf{e}_t; 0, \mathscr{S}_t) . && (5.19)
\end{aligned}$$

$$(M_t^f(j), W_t^f(i,j)) = ForwardSwitchFilter(M_{t-1}^f, Z, L_t)$$

$$\begin{aligned}
\alpha_t^f(i,j) &= M_{t-1}^f(i)L_t^f(i,j)Z(i,j) , && (5.20)\\
a_t^f(j) &= \sum_{i'} \alpha_t^f(i',j) , && (5.21)\\
M_t^f(j) &= \frac{a_t^f(j)}{\sum_{j'} a_t^f(j')} , && (5.22)\\
W_t^f(i,j) &= \frac{\alpha_t^f(i,j)}{\sum_{j'} a_t^f(j')} / M_t^f(j). && (5.23)
\end{aligned}$$

$$(\mu_{t|t}(j), V_{t|t}(j)) = Collapse(\tilde{\mu}_{t|t}, \tilde{V}_{t|t}, W_t^f)$$

$$\begin{aligned}
\mu_{t|t}(j) &= \sum_i W_t^f(i,j)\tilde{\mu}_{t|t}(i,j) && (5.24)\\
V_{t|t}(j) &= \sum_i W_t^f(i,j)(\tilde{V}_{t|t}(i,j) + (\tilde{\mu}_{t|t}(i,j) - \mu_{t|t}(j))(\tilde{\mu}_{t|t}(i,j) - \mu_{t|t}(j))^{\mathsf{T}}) , && (5.25)
\end{aligned}$$

For $t = 1, \cdots, T$. Note, $M_t^f(i) = Prob(S_t = i | y_{1:t})$, with the initial condition $M_0^f = \pi$.

The analytical partial derivatives of the three operators above are presented in Appendix A, sections 5.9.1, 5.9.2, and 5.9.3.

5.4.2 Smoothing Step

The smoothing step of the SKF algorithm for the switching variables takes the following form [42]:

$(M_t^s(i)) = BackwardSwitchSmooth(M_t^f,\ M_{t+1}^s,\ Z)$

$$a_t^s(i,j) = M_t^f(i)Z(i,j)\ , \tag{5.26}$$

$$b_t^s(i,j) = \frac{a_t^s(i,j)}{\sum_{i'} a_t^s(i',j)}\ , \tag{5.27}$$

$$M_t^s(i) = \sum_{j'} b_t^s(i,j')M_{t+1}^s(j')\ , \tag{5.28}$$

for $t = T-1,\cdots,1$. Note, $M_t^s(i) = Prob(S_t = i|y_{1:T})$, with the initial condition $M_T^s = M_T^f$.

5.5 Error Gradient Calculations

We start from the filtering step of the SKF algorithm and calculate the analytical partial derivatives of each node output(s) with respect to its input(s), as we move forward in time. Next, smoothing of the switching variables is performed and the corresponding analytical gradients are calculated. The back-propagation algorithm starts from the reverse direction (from the output of the smoothed switching variables or the mode proportions) and propagates the gradient information backward (starting from time T) through the smoothed switching variables, and finally the filtered variables (end ing in time t=1).

5.5.1 Error Gradient with Respect to Smoothed Switching Variables

Derivatives of the error with respect to the mode proportions η_i are given by Eq. (5.6) (in the case of binary outcomes) and Eq. (5.10) (in the case of multinomial outcomes). Next, the error is backpropagated through the smoothed switching variables, as follows:

$$\begin{aligned}
\frac{\partial E}{\partial M_1^s(i)} &= \frac{1}{T}\frac{\partial E}{\partial \eta_i}\ , \\
\frac{\partial E}{\partial M_t^s(i)} &= \frac{1}{T}\frac{\partial E}{\partial \eta_i} + \sum_{j'} \frac{\partial E}{\partial M_{t-1}^s(j')} b_{t-1}^s(j',i)\ ,\ t = 2\cdots T
\end{aligned} \tag{5.29}$$

We also compute the following derivatives:

$$\frac{\partial E}{\partial a_t^s(i,j)} = \sum_{k'} \frac{\partial E}{\partial M_t^s(k')} M_{t+1}^s(j)\left[\frac{\delta_{k',i}}{\sum_{i'} a_t^s(i',j)} - \frac{a_t^s(k',j)}{(\sum_{i'} a_t^s(i',j))^2}\right]\ ,\ t = 1\cdots T-1 \tag{5.30}$$

5.5.2 Error Gradient with Respect to Filtered Switching Variables

Next, derivatives of the error with respect to the filtered switching variables can be calculated as follows:

$$\begin{aligned}
\frac{\partial E}{\partial M_T^f(i)} &= \frac{\partial E}{\partial M_T^s(i)}\,, \\
\frac{\partial E}{\partial M_t^f(i)} &= \sum_{j'} \frac{\partial E}{\partial a_t^s(i,j')} Z(i,j') + \frac{\partial E}{\partial W_{t+1}^f} \odot \frac{\partial W_{t+1}^f}{\partial M_t^f(i)} + \frac{\partial E}{\partial M_{t+1}^f} \odot \frac{\partial M_{t+1}^f}{\partial M_t^f(i)}\,, \; t = T-1 \cdots 1
\end{aligned} \tag{5.31}$$

where

$$\begin{aligned}
\frac{\partial E}{\partial W_T^f(i,j)} &= 0, \forall\, i,j\,, \\
\frac{\partial E}{\partial W_t^f(i,j)} &= \frac{\partial E}{\partial \mu_t^f(j)} \odot \frac{\partial \mu_t^f(j)}{\partial W_t^f(i,j)} + \frac{\partial E}{\partial V_t^f(j)} \odot \frac{\partial V_t^f(j)}{\partial W_t^f(i,j)}\,, \; t = T-1 \cdots 1\,.
\end{aligned} \tag{5.32}$$

5.5.3 Error Gradient with Respect to Filtered State Variables

Although the marginal distributions of the state variable are Gaussian random variables, during the error back-propagation they can be treated as real-valued means (vectors) and covariances (matrices). The following gradients of the error with respect to the state variables are recursively calculated:

$$\begin{aligned}
\frac{\partial E}{\partial \mu_t^f(i)[m]} &= \sum_{j'} \frac{\partial E}{\partial L_{t+1}^f(i,j')} \frac{\partial L_{t+1}^f(i,j')}{\partial \mu_t^f(i)[m]} + \sum_{j'} \frac{\partial E}{\partial \tilde{\mu}_{t+1}^f(i,j')} \odot \frac{\partial \tilde{\mu}_{t+1}^f(i,j')}{\partial \mu_t^f(i)[m]}\,, \; t = 1 \cdots T-2\,, \\
\frac{\partial E}{\partial \mu_t^f(i)[m]} &= \sum_{j'} \frac{\partial E}{\partial L_{t+1}^f(i,j')} \frac{\partial L_{t+1}^f(i,j')}{\partial \mu_t^f(i)[m]}\,, \; t = T-1\,,
\end{aligned} \tag{5.33}$$

$$\begin{aligned}
\frac{\partial E}{\partial V_t^f(i)[m,n]} &= \sum_{j'} \frac{\partial E}{\partial L_{t+1}^f(i,j')} \frac{\partial L_{t+1}^f(i,j')}{\partial V_t^f(i)[m,n]} + \sum_{j'} \frac{\partial E}{\partial \tilde{V}_{t+1}^f(i,j')} \odot \frac{\partial \tilde{V}_{t+1}^f(i,j')}{\partial V_t^f(i)[m,n]} \\
&\quad + \sum_{j'} \frac{\partial E}{\partial \tilde{\mu}_{t+1}^f(i,j')} \odot \frac{\partial \tilde{\mu}_{t+1}^f(i,j')}{\partial V_t^f(i)[m,n]}\,, \; t = 1 \cdots T-2\,, \\
\frac{\partial E}{\partial V_t^f(i)[m,n]} &= \sum_{j'} \frac{\partial E}{\partial L_{t+1}^f(i,j')} \frac{\partial L_{t+1}^f(i,j')}{\partial V_t^f(i)[m,n]}\,, \; t = T-1\,,
\end{aligned} \tag{5.34}$$

where

$$\frac{\partial E}{\partial L_t^f(i,j)} = \frac{\partial E}{\partial W_t^f} \odot \frac{\partial W_t^f}{\partial L_t^f(i,j)} + \frac{\partial E}{\partial M_t^f} \odot \frac{\partial M_t^f}{\partial L_t^f(i,j)}\,, \; t = 1 \cdots T\,, \tag{5.35}$$

and

$$\frac{\partial E}{\partial \tilde{\mu}_t^f(i,j)[m]} = \frac{\partial E}{\partial \mu_t^f(j)} \odot \frac{\partial \mu_t^f(j)}{\partial \tilde{\mu}_t^f(i,j)[m]} + \frac{\partial E}{\partial V_t^f(j)} \odot \frac{\partial V_t^f(j)}{\partial \tilde{\mu}_t^f(i,j)[m]}, \; t = 1 \cdots T-1, \tag{5.36}$$

$$\frac{\partial E}{\partial \tilde{V}_t^f(i,j)[m]} = \frac{\partial E}{\partial V_t^f(j)} \odot \frac{\partial V_t^f(j)}{\partial \tilde{V}_t^f(i,j)[m,n]}, \; t = 1 \cdots T-1. \tag{5.37}$$

Note that, $\partial \tilde{V}_{t+1}^f(i,j')/\partial \mu_t^f(i) = 0$ for all i, j, and $\partial \mu_t^f(j)/\partial \tilde{V}_t^f(i,j) = 0$ for all i, j.

5.5.4 Error Gradient with Respect to Model Parameters

We finally arrive at the error gradients with respect to the model parameters. The derivatives with respect to the Markov switching state transition matrix are given by:

$$\frac{\partial E}{\partial Z(i,j)} = \sum_{t=1}^{T} \frac{\partial E}{\partial M_t^f(j)} \odot \frac{\partial M_t^f(j)}{\partial Z(i,j)} + \sum_{t=1}^{T-1} \frac{\partial E}{\partial a_t^s(i,j)} M_t^f(i) \tag{5.38}$$

For the other model parameters the error gradients are as follows:

$$\begin{aligned}\frac{\partial E}{\partial A(j)[m,n]} &= \sum_{i'}[\sum_{t=1}^{T-1} \frac{\partial E}{\partial \tilde{\mu}_t^f(i',j)} \odot \frac{\partial \tilde{\mu}_t^f(i',j)}{\partial A(j)[m,n]} \\ &+ \sum_{t=1}^{T-1} \frac{\partial E}{\partial \tilde{V}_t^f(i',j)} \odot \frac{\partial \tilde{V}_t^f(i',j)}{\partial A(j)[m,n]} + \sum_{t=1}^{T} \frac{\partial E}{\partial L_t^f(i',j)} \frac{\partial L_t^f(i',j)}{\partial A(j)[m,n]}].\end{aligned} \tag{5.39}$$

$$\begin{aligned}\frac{\partial E}{\partial C(j)[m,n]} &= \sum_{i'}[\sum_{t=1}^{T-1} \frac{\partial E}{\partial \tilde{\mu}_t^f(i',j)} \odot \frac{\partial \tilde{\mu}_t^f(i',j)}{\partial C(j)[m,n]} \\ &+ \sum_{t=1}^{T-1} \frac{\partial E}{\partial \tilde{V}_t^f(i',j)} \odot \frac{\partial \tilde{V}_t^f(i',j)}{\partial C(j)[m,n]} + \sum_{t=1}^{T} \frac{\partial E}{\partial L_t^f(i',j)} \frac{\partial L_t^f(i',j)}{\partial C(j)[m,n]}].\end{aligned} \tag{5.40}$$

$$\begin{aligned}\frac{\partial E}{\partial Q(j)[m,n]} &= \sum_{i}[\sum_{t=1}^{T-1} \frac{\partial E}{\partial \tilde{\mu}_t^f(i',j)} \odot \frac{\partial \tilde{\mu}_t^f(i',j)}{\partial Q(j)[m,n]} \\ &+ \sum_{t=1}^{T-1} \frac{\partial E}{\partial \tilde{V}_t^f(i',j)} \odot \frac{\partial \tilde{V}_t^f(i',j)}{\partial Q(j)[m,n]} + \sum_{t=1}^{T} \frac{\partial E}{\partial L_t^f(i',j)} \frac{\partial L_t^f(i',j)}{\partial Q(j)[m,n]}].\end{aligned} \tag{5.41}$$

$$\begin{aligned}\frac{\partial E}{\partial R(j)[m,n]} &= \sum_{i'}[\sum_{t=1}^{T-1} \frac{\partial E}{\partial \tilde{\mu}_t^f(i',j)} \odot \frac{\partial \tilde{\mu}_t^f(i',j)}{\partial R(j)[m,n]} \\ &+ \sum_{t=1}^{T-1} \frac{\partial E}{\partial \tilde{V}_t^f(i',j)} \odot \frac{\partial \tilde{V}_t^f(i',j)}{\partial R(j)[m,n]} + \sum_{t=1}^{T} \frac{\partial E}{\partial L_t^f(i',j)} \frac{\partial L_t^f(i',j)}{\partial R(j)[m,n]}].\end{aligned} \tag{5.42}$$

5.6 Optimization

Our objective is to minimize the cost function in Eqs. (5.3) or (5.8), subject to the constraints that all Q and R covariance matrices must remain positive-definite, and elements of Z must be positive

and add up to one. This problem is equivalent to solving the unconstrained optimization problem:

$$\text{Minimize } E_{\{A(j),\, \bar{Q}(j),\, C(j),\, \bar{R}(j),\, \bar{Z},\beta\},j=1\cdots J}\,, \tag{5.43}$$

where the β are the regression coefficients, $Q(j) = exp(\bar{Q}(j))$ and $R(j) = exp(\bar{R}(j))$ (both assumed diagonal for simplicity), with the following gradients:

$$\frac{\partial E}{\partial \bar{Q}(j)[m,n]} = \bar{Q}(j)[m,n]\frac{\partial E}{\partial Q(j)[m,n]}, \tag{5.44}$$

$$\frac{\partial E}{\partial \bar{R}(j)[m,n]} = \bar{R}(j)[m,n]\frac{\partial E}{\partial R(j)[m,n]}. \tag{5.45}$$

Moreover, $Z(i,j) = \frac{exp(\bar{Z}(i,j))}{\sum_{j'} exp(\bar{Z}(i,j'))}$, which results in the following gradient for $\bar{Z}$:

$$\frac{\partial E}{\partial \bar{Z}(i,j)} = \frac{\partial E}{\partial \bar{Z}} \odot \frac{\partial Z}{\partial \bar{Z}(i,j)}, \tag{5.46}$$

where

$$\frac{\partial Z(k,l)}{\partial \bar{Z}(i,j)} = \delta_{i,k} Z(i,j)(\delta_{j,l} - Z(k,l)). \tag{5.47}$$

5.6.1 EM-based Initialization

Let us assuming a binary outcome classification problem, with diagonal matrices Q, R, and a constant C matrix for an MVAR model of order P. The optimization problem in Eq. (5.43) will have $J \times (M^2 \times P + M + M) + J^2 + J$ free parameters, where J is the number of modes. As an example, for $J = 4$, $P = 2$ and $M = 2$ (a second order bivariate AR) we will have 68 parameters to adjust. In the case of the MIMIC II example in the previous section ($J = 20$, $P = 3$, $M = 1$) the number of free parameters is 520.

The large number of free parameters prohibits the use of standard optimization techniques, due to overfitting, convergence issues, and stability of the Kalman filter. In practice, we initialize the switching Kalman filter using a few iterations of EM (10-15 iterations), and then switch to a gradient descent-based optimizer such as the LBFGS method (for Limited-Memory low-rank Hessian approximation). The efficiency of the LBFGS technique is due to using an approximation of the Hessian (second order partial derivatives required by the Newton's optimization method) by accumulating information contained in the gradient and error function values during optimization. A Matlab implementation of the LBFGS algorithm can be found in the *minFunc* optimization pack-

aged (available at: `http://www.di.ens.fr/~mschmidt/Software/minFunc.html`).

The intuition behind the EM-based initialization step is as follows. We are assuming that all the information needed for making a prediction (or classification) is within certain dynamical features of the time-series. However, the optimization surface may include many local minima, and therefore is sensitive to initialization. By maximizing the expected log-likelihood of data (with no labels), EM attempts to learn the different dynamical features that are present within the time-series. However, certain features (such as artifacts) may not be relevant to the outcomes of interest. Nevertheless, early stopping of the EM mitigates the problem of learning uninformative features through overfitting, and puts the parameter values in the appropriate place (within the parameter space) for further supervised learning. At this point, outcome-discriminative (or supervised) learning can be used to allocate more resources (in terms of modes and their associated eigenvalues or poles) to predictive features of the time-series.

5.6.2 Notes on Implementation

Note that both the E-step of the EM algorithm and the gradient computation-step of the supervised learning can be done in parallel for each time-series. Therefore, in problems involving hundreds to thousands of time-series, one may run these steps on a computer cluster, for instance using Matlab's Parallel Computing toolbox.

5.7 Some Illustrative Examples

5.7.1 Simulated Time-Series with Multinomial Outcomes

We simulated 200 bivariate ($M = 2$) AR time-series of model order $P = 2$, consisting of $J = 4$ different modes (for the detail of the simulation see Appendix 5.10.1). The time-series were divided into four categories, according to the proportion of time spent within each mode, with the mode transition probabilities as shown in Fig. 5-3. Additionally, when generating each time-series, the dynamics matrices $A_1, \cdots, A_4$ were individually perturbed by adding a white Gaussian noise of standard deviation of 0.05 to each of their elements; this in effect creates some dissimilarities between the modes across the different time-series. Additionally, two 10-sample-long artifactual segments (uniform random noise in the interval of [0, 15]) were inserted at random locations within each time-series. Four sample time-series from each of the four categories are shown in Fig. 5-1.

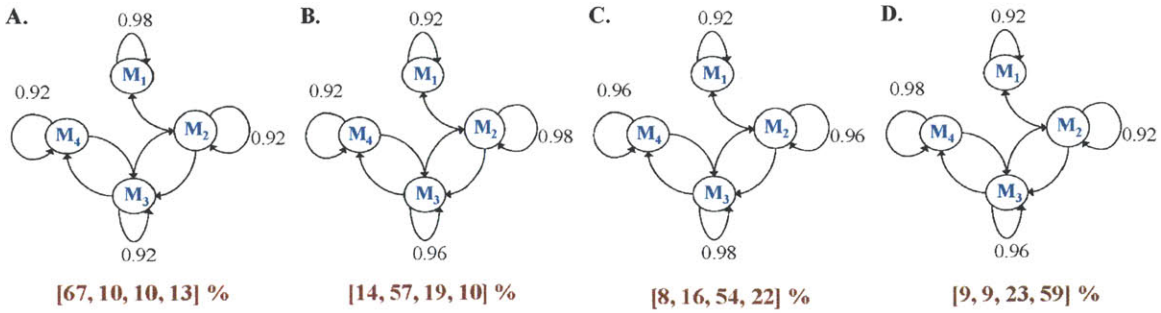

Figure 5-3: Transition diagram for the four categories
The time-series were divided into 4 categories (A, B, C, D), each statistically having different proportions of four modes ($M_1, \cdots, M_4$). Specifically, the stationary distribution of the four categories were [0.67, 0.10, 0.10, 0.13], [0.14, 0.57, 0.19, 0.10], [0.08, 0.16, 0.54, 0.22], and [0.09, 0.09 , 0.23, 0.59]. Note that only the important transition links are shown; the actual transition matrices were fully connected.

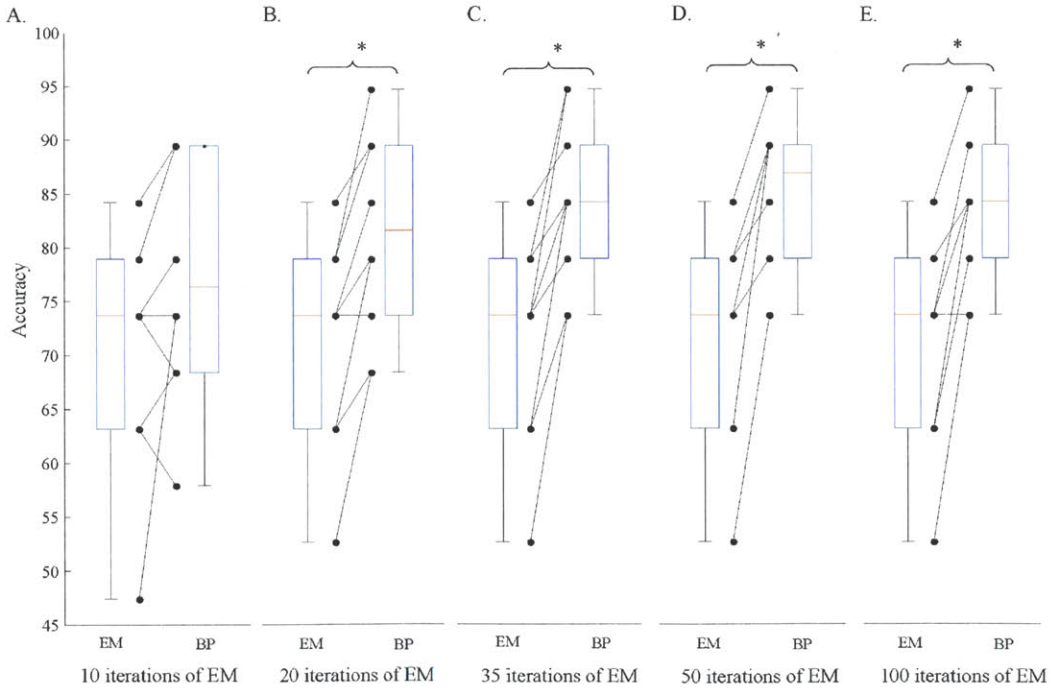

Figure 5-4: Comparison of EM and supervised-learning.
Panels A-C show the influence of EM initialization on the performance of the supervised learning (BP). In each case, 30 iterations of the LBFGS optimizer followed the EM runs. The presented results are based on 10 fold cross-validated classification performance (represented with closed circles); each boxplot is constructed using the classification performance on 10 testing folds.

Fig. 5-4 (panels A-E) summarizes the simulation study results. The classification performance using EM alone slightly improves between 10 and 20 iterations; however, further iterations of the EM do not result in improved performance on the testing sets. Notably, the figure demonstrates the dependence of the proposed outcome-discriminative (or supervised) learning on the EM initialization step. With only 10 iterations of EM (panel A), occasionally the supervised learning algorithm results in a lower performance on the testing data. However, after 20 iterations of EM (panel B, EM median accuracy = $\sim$74%) the supervised learning produces consistent improvement in performance (BP median accuracy = $\sim$83%, $p<0.05$; signed-rank test). Additional EM steps, up to 50 iterations, marginally improve the supervised learning performance (BP median accuracy = $\sim$87). However, higher iterations of the EM seems to lower the performance of the supervised learning (presumably due to overfitting of artifacts, making it harder may be harder for the optimizer to escape the corresponding local minima).

5.7.2 Multinomial Decoding of Posture: Tilt-Table experiment

Let us return to the tilt-table example of Chapter 4, Section 4.3.2. Here we construct a substantially more difficult multinomial classification task involving the four maneuvers depicted in Fig. 4-3, that is, supine, slow tilt, rapid tilt, and standing up. This task can be viewed as a sequence labeling problem, where variable-length segments of a time-series need to be labeled. Such sequence labeling problems are common in speech recognition.

The results presented in Fig. 5-5 are based on 10 fold cross-validated performance on the testing data. Similar to Chapter 4, we used four modes, each corresponding to an AR model of order three, to model the bivariate time-series of heart rate and blood pressure. We initialized the supervised learning algorithm using 50 iterations of EM, followed by 30 iterations of supervised learning.

Figure 5-5 panel A shows a comparison of the EM-based classification vs. supervised learning. The EM-based classification achieves an accuracy of 80.8 [71.1 84.6]; however, the performance is quite varied across the different testing folds. In comparison, application of supervised learning yields a median accuracy of 84.6 [84.6 90.4] with smaller interquartile range.

Figure 5-5 panel B shows the multinomial probability of the observed outcomes, given the learned model parameters, before and after supervised learning. Note that this is the quantity we try to maximize by supervised learning, although on the training data. Notably, where the probability of the outcomes under the model is the worst (bottom three folds), the application of supervised

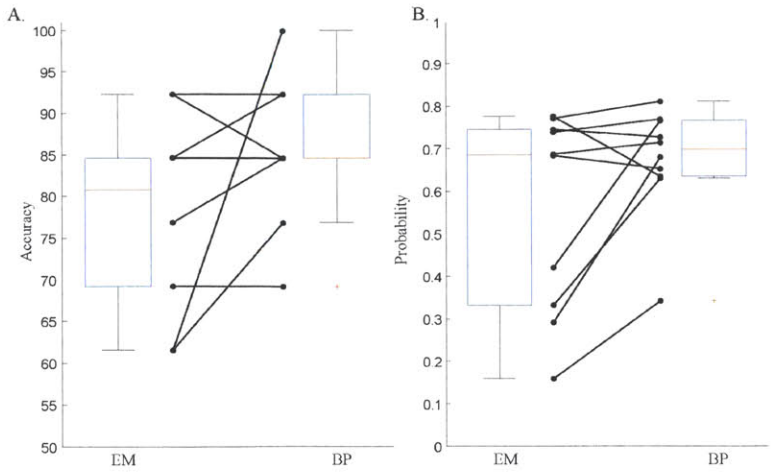

Figure 5-5: 10 fold cross-validation results.
Each closed-circle represents the classification performance on the testing data for one out of the 10 folds. Panel A shows a comparison of the accuracy achieved using EM vs. BP on the testing data: EM Accuracy 80.8 [71.1 84.6] percent (median [Interquartile]) vs. BP Accuracy 84.6 [84.6 90.4] percent. Panel B shows a comparison of the two techniques in terms of multinomial probability of observed outcomes (on the testing data) based on the model: EM testing prob 0.69 [0.35 0.74] vs. BP testing prob 0.70 [0.64 0.76]. These results indicate that the supervised algorithm is especially effective where the EM performance is poor.

learning yields the largest improvement in performance.

Figure 5-6 presents a comparison of the two techniques in terms of the individual classification tasks. As shown in panel A, the EM-based classification incurs the largest error in distinguishing between the fast tilt and standing up. Panel B shows that the supervised learning improves the distinction between fast tilting and standing-up. Notably, in almost all cases it classifies the standing-up category correctly, and yields a small improvement on classification of slow-tilt events.

Physiological Interpretation of the Discovered Dynamics

As extensively discussed in Chapter 2, the AR coefficients corresponding to each of the dynamical modes can be used to untangle the directional relationships among the modeled variables. Additionally, useful descriptive indices of the individual time-series can be obtained from the dynamical modes. For instance, the ratio of the low-frequency power (LF: corresponding to periods of 6-20 beats, mainly a function of sympathetic activity) to high-frequency power (HF: corresponding to periods of 2-5 beats, mainly reflecting parasympathetic or vagal activation) of the RR interval time-series is considered to mirror sympathovagal balance or to reflect the sympathetic modulations. The results presented in Fig. 5-7 show a reduction in HF band power during each of the three tilting

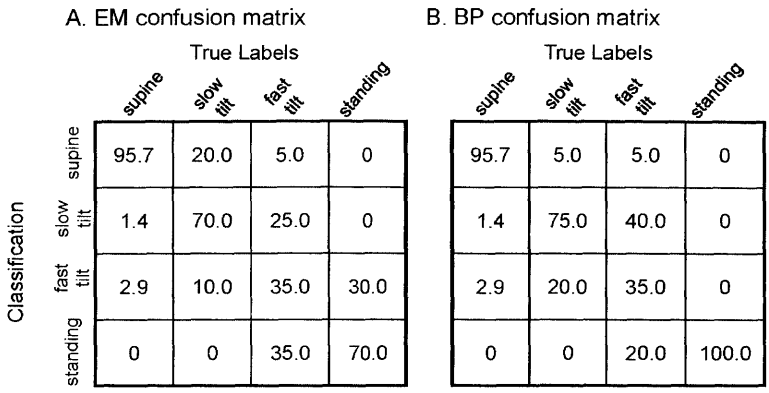

Figure 5-6: Average confusion matrices.
Panel A shows the average confusion matrix for the EM-based classification. Panel B shows the average confusion matrix for the supervised classification. Accuracy is defined as the sum of the diagonal elements of the confusion matrix divided by the sum total of all the elements of the matrix.

events, which results in an increased LF/HF ratio, indicating increased sympathetic modulations. These results were obtained by (1) calculating the power spectrum of each mode, and (2) calculating a weighted average of the spectra within each segment, where the weights were given by the mode proportions.

5.8 Discussion

We presented a technique for outcome-discriminative learning of dynamics within a cohort time-series. The main idea of our approach was to present the learning algorithm with the outcomes (or labels) corresponding to each time-series, in order to learn time-series features that are most relevant to the discriminative task of distinguishing among the labels. The utility of the technique was demonstrated using both simulated data and time-series recordings from a human laboratory tilt-table study.

Using simulated time-series, we showed that the supervised learning algorithm provides a significant improvement over EM, given proper EM-based initialization. Notably, the proposed EM-based initialization step is qualitatively similar to the unsupervised learning step used for training Deep Belief Networks (DBN) [111], where the unsupervised learning step is known to significantly improve the predictive performance of these algorithms [112].

Although we developed our outcome-discriminative learning algorithm in the context of switching linear dynamical systems (SLDS), one can apply a similar technique to the problem of learning

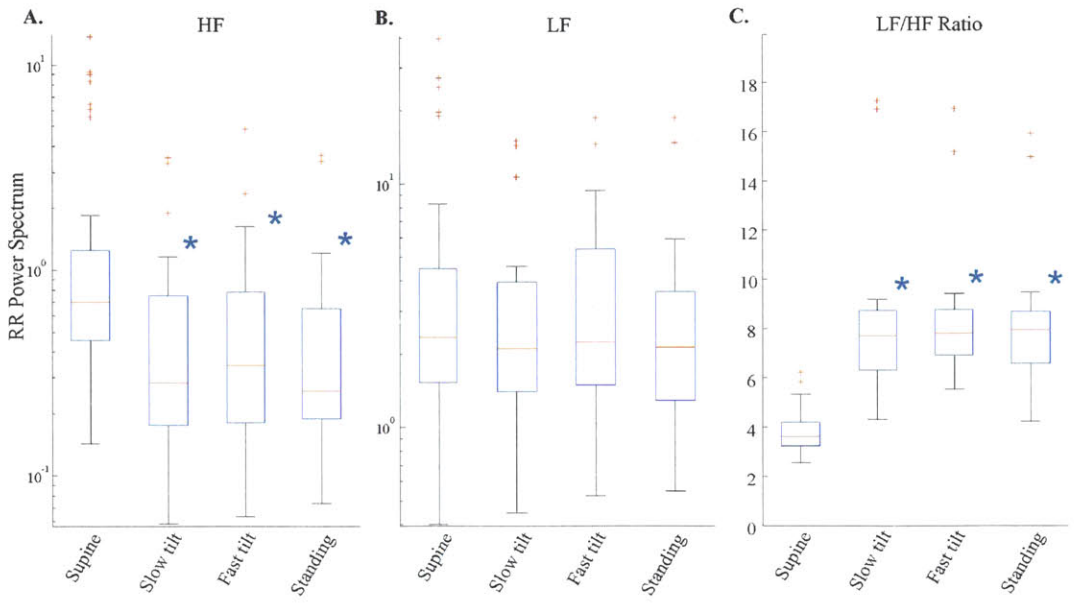

Figure 5-7: Physiological interpretation of learned dynamics
Average parametric power of the RR intervals with the HF (period of 2-5 beats) and LF (periods of 6-20 beats) bands are shown in panels A and B, respectively. Panel C shows the LF/HF ratio. The $\star$ symbol indicates a significant change from baseline ($p<0.05$; Kruskal-Wallis nonparametric ANOVA test).

marginal distributions and model parameters in a variety of other probabilistic dynamic Bayesian networks (defined on directed acyclic graphs, including the SLDS as special case). Learning could be done using an equivalent neural network representation to directly optimize a cost function on the inferred marginals.

The approach presented here has a similar flavor as the works of Memisevic (2006) [113], Eaton (2009) [114], Stoyanov (2011) [115] and Domke (2012) [116] in the context of learning and inference in Markov random fields. All the previous works in this area have focused on *discrete* models, and used direct optimization to minimize an empirical risk (or cost) function of the inferred marginal distributions.

As demonstrated through the tilt-table example, the proposed approach has the added advantage of having physiological interpretability. Since the features used for prediction are based on the dynamics of the underlying time-series, one can link back the most predictive features for a given outcome of interest to the underlying physiology. For instance, tilting is known to disrupt the sympathovagal balance in the direction of increased sympathetic activation. Notably, modes that were most predictive of the tilting events had a larger LF/HF ratio, indicating increased sympathetic modulations.

Given the sheer volume of multivariate time-series recorded in modern clinical databases, one of the advantages of the proposed technique is its scalability, due to the parallel implementation of the inference step of the SKF algorithm (and thus the gradient calculations). Therefore, the learning of the dynamical modes, and the Markov transition matrix, can be performed efficiently by exploiting the modern multi-core processing units (CPU cores) and cluster computation.

5.9 Appendix A

5.9.1 Analytical Derivatives of the Kalman Filter

Recall the Kalman filter operator is defined as:

$$(\tilde{\mu}_{t|t}(i,j),\tilde{V}_{t|t}(i,j),L_t(i,j)) = KalmanFilter(\mu_{t-1|t-1}(i),V_{t-1|t-1}(i),\mathbf{y}_t(i);A(j),C(j),Q(j),R(j)) .$$

The analytical derivatives of the Kalman filter operator is given by

$$\frac{\partial \tilde{\mu}_{t|t-1}}{\partial A[m,n]} = \delta_{m,n}\mu_{t-1|t-1} \tag{5.48}$$

$$\frac{\partial \tilde{\mu}_{t|t-1}}{\partial \mu_{t-1|t-1}(m)} = A\delta_m \tag{5.49}$$

$$\frac{\partial \tilde{V}_{t|t-1}}{\partial A[m,n]} = \delta_{m,n}V_{t-1|t-1}A^{\mathsf{T}} + AV_{t-1|t-1}\delta_{m,n}^{\mathsf{T}} \tag{5.50}$$

$$\frac{\partial \tilde{V}_{t|t-1}}{\partial Q[m,n]} = \delta_{m,n}^{n,m} \tag{5.51}$$

$$\frac{\partial \tilde{V}_{t|t-1}}{\partial V_{t-1|t-1}[m,n]} = A\delta_{m,n}^{n,m}A^{\mathsf{T}} \tag{5.52}$$

$$\frac{\partial \mathbf{e}_t}{\partial A[m,n]} = -C\delta_{m,n}\mu_{t-1|t-1} \tag{5.53}$$

$$\frac{\partial \mathbf{e}_t}{\partial C[m,n]} = -\delta_{m,n}A\mu_{t-1|t-1} \tag{5.54}$$

$$\frac{\partial \mathbf{e}_t}{\partial \mu_{t-1|t-1}[m]} = -CA\delta_m \tag{5.55}$$

$$\frac{\partial \mathscr{S}_t}{\partial A[m,n]} = C\frac{\partial \tilde{V}_{t|t-1}}{\partial A[m,n]}C^{\mathsf{T}} \tag{5.56}$$

$$\frac{\partial \mathscr{S}_t}{\partial C[m,n]} = \delta_{m,n}\tilde{V}_{t|t-1}C^{\mathsf{T}} + C\tilde{V}_{t|t-1}\delta_{m,n}^{\mathsf{T}} \tag{5.57}$$

$$\frac{\partial \mathscr{S}_t}{\partial Q[m,n]} = C\frac{\partial \tilde{V}_{t|t-1}}{\partial Q[m,n]}C^{\mathsf{T}} \tag{5.58}$$

$$\frac{\partial \mathscr{S}_t}{\partial R[m,n]} = \delta_{m,n}^{n,m} \tag{5.59}$$

$$\frac{\partial \mathscr{S}_t}{\partial V_{t-1|t-1}[m,n]} = C\frac{\partial \tilde{V}_{t|t-1}}{\partial V_{t-1|t-1}[m,n]}C^{\mathsf{T}} \tag{5.60}$$

$$\frac{\partial G_t}{\partial A[m,n]} = (\frac{\partial \tilde{V}_{t|t-1}}{\partial A[m,n]}C^{\mathsf{T}} - G_t\frac{\partial \mathscr{S}_t}{\partial A[m,n]})\mathscr{S}_t^{-1} \tag{5.61}$$

$$\frac{\partial G_t}{\partial C[m,n]} = (\tilde{V}_{t|t-1}\delta_{m,n}^{\mathsf{T}} - G_t\frac{\partial \mathscr{S}_t}{\partial C[m,n]})\mathscr{S}_t^{-1} \tag{5.62}$$

$$\frac{\partial G_t}{\partial Q[m,n]} = (\frac{\partial \tilde{V}_{t|t-1}}{\partial Q[m,n]}C^{\mathsf{T}} - G_t\frac{\partial \mathscr{S}_t}{\partial Q[m,n]})\mathscr{S}_t^{-1} \tag{5.63}$$

$$\frac{\partial G_t}{\partial R[m,n]} = (-G_t\frac{\partial \mathscr{S}_t}{\partial R[m,n]})\mathscr{S}_t^{-1} \tag{5.64}$$

$$\frac{\partial G_t}{\partial V_{t-1|t-1}[m,n]} = (\frac{\partial \tilde{V}_{t|t-1}}{\partial V_{t-1|t-1}[m,n]}C^{\mathsf{T}} - G_t\frac{\partial \mathscr{S}_t}{\partial V_{t-1|t-1}[m,n]})\mathscr{S}_t^{-1} \tag{5.65}$$

$$\frac{\partial \tilde{\mu}_{t|t}}{\partial A[m,n]} = \frac{\partial \tilde{\mu}_{t|t-1}}{\partial A[m,n]} + \frac{\partial G_t}{\partial A[m,n]}\mathbf{e}_t + G_t\frac{\partial \mathbf{e}_t}{\partial A[m,n]} \tag{5.66}$$

$$\frac{\partial \tilde{\mu}_{t|t}}{\partial C[m,n]} = \frac{\partial G_t}{\partial C[m,n]}\mathbf{e}_t + G_t\frac{\partial \mathbf{e}_t}{\partial C[m,n]} \tag{5.67}$$

$$\frac{\partial \tilde{\mu}_{t|t}}{\partial Q[m,n]} = \frac{\partial G_t}{\partial Q[m,n]}\mathbf{e}_t \tag{5.68}$$

$$\frac{\partial \tilde{\mu}_{t|t}}{\partial R[m,n]} = \frac{\partial G_t}{\partial R[m,n]}\mathbf{e}_t \tag{5.69}$$

$$\frac{\partial \tilde{\mu}_{t|t}}{\partial \mu_{t-1|t-1}[m]} = \frac{\partial \tilde{\mu}_{t|t-1}}{\partial \mu_{t-1|t-1}(m)} + G_t\frac{\partial \mathbf{e}_t}{\partial \mu_{t-1|t-1}[m]} \tag{5.70}$$

$$\frac{\partial \tilde{\mu}_{t|t}}{\partial V_{t-1|t-1}[m,n]} = \frac{\partial G_t}{\partial V_{t-1|t-1}[m,n]}\mathbf{e}_t \tag{5.71}$$

$$\frac{\partial \tilde{V}_{t|t}}{\partial A[m,n]} = -\frac{\partial G_t}{\partial A[m,n]} C\tilde{V}_{t|t-1} + (I - G_t C)\frac{\partial \tilde{V}_{t|t-1}}{\partial A[m,n]} \tag{5.72}$$

$$\frac{\partial \tilde{V}_{t|t}}{\partial C[m,n]} = -(\frac{\partial G_t}{\partial C[m,n]} C + G_t \delta_{m,n})\tilde{V}_{t|t-1} \tag{5.73}$$

$$\frac{\partial \tilde{V}_{t|t}}{\partial Q[m,n]} = -\frac{\partial G_t}{\partial Q[m,n]} C\tilde{V}_{t|t-1} + (I - G_t C)\frac{\partial \tilde{V}_{t|t-1}}{\partial Q[m,n]} \tag{5.74}$$

$$\frac{\partial \tilde{V}_{t|t}}{\partial R[m,n]} = -\frac{\partial G_t}{\partial R[m,n]} C\tilde{V}_{t|t-1} \tag{5.75}$$

$$\frac{\partial \tilde{V}_{t|t}}{\partial V_{t-1|t-1}[m,n]} = -\frac{\partial G_t}{\partial V_{t-1|t-1}[m,n]} C\tilde{V}_{t|t-1} + (I - G_t C)\frac{\partial \tilde{V}_{t|t-1}}{\partial V_{t-1|t-1}[m,n]} \tag{5.76}$$

$$\frac{\partial L_t}{\partial \mathbf{e}_t} = -L_t^f \mathscr{S}_t^{-1} \mathbf{e}_t \tag{5.77}$$

$$\frac{\partial L_t}{\partial \mathscr{S}_t} = -\frac{1}{2} L_t^f (\mathscr{S}_t^{-1} - \mathscr{S}_t^{-1} \mathbf{e}_t \mathbf{e}_t^\top \mathscr{S}_t^{-1}) \tag{5.78}$$

$$\frac{\partial L_t}{\partial A[m,n]} = \frac{\partial L_t}{\partial \mathbf{e}_t} \odot \frac{\partial \mathbf{e}_t}{\partial A[m,n]} + \frac{\partial L_t}{\partial \mathscr{S}_t} \odot \frac{\partial \mathscr{S}_t}{\partial A[m,n]} \tag{5.79}$$

$$\frac{\partial L_t}{\partial C[m,n]} = \frac{\partial L_t}{\partial \mathbf{e}_t} \odot \frac{\partial \mathbf{e}_t}{\partial C[m,n]} + \frac{\partial L_t}{\partial \mathscr{S}_t} \odot \frac{\partial \mathscr{S}_t}{\partial C[m,n]} \tag{5.80}$$

$$\frac{\partial L_t}{\partial Q[m,n]} = \frac{\partial L_t}{\partial \mathscr{S}_t} \odot \frac{\partial \mathscr{S}_t}{\partial Q[m,n]} \tag{5.81}$$

$$\frac{\partial L_t}{\partial R[m,n]} = \frac{\partial L_t}{\partial \mathscr{S}_t} \odot \frac{\partial \mathscr{S}_t}{\partial R[m,n]} \tag{5.82}$$

$$\frac{\partial L_t}{\partial \mu_{t-1|t-1}[m]} = \frac{\partial L_t}{\partial \mathbf{e}_t} \odot \frac{\partial \mathbf{e}_t}{\partial \mu_{t-1|t-1}[m]} \tag{5.83}$$

$$\frac{\partial L_t}{\partial V_{t-1|t-1}[m,n]} = \frac{\partial L_t}{\partial \mathscr{S}_t} \odot \frac{\partial \mathscr{S}_t}{\partial V_{t-1|t-1}[m,n]} \tag{5.84}$$

5.9.2 Analytical Derivatives of the Filtered Switching Variables

The analytic derivatives of the forward filtering operator:
$(M_t^f(j), W_t^f(i,j)) = ForwardSwitchFilter(M_{t-1}^f, Z, L_t)$, is given as follows:

$$\frac{\partial M_t^f(i)}{\partial M_{t-1}^f(j)} = \sum_{k'} L_t^f(j,i) Z(j,i) [\frac{\delta_{k',j}}{\sum_{i'} a_t^f(i')} - \frac{a_t^f(i)}{(\sum_{i'} a_t^f(i'))^2}] \ , \ t = 1 \cdots T. \tag{5.85}$$

$$\frac{\partial M_t^f(k)}{\partial Z(i,j)} = M_{t-1}^f(i) L_t^f(i,j) [\frac{\delta_{k',j}}{\sum_{k'} a_t^f(k')} - \frac{a_t^f(k)}{(\sum_{k'} a_t^f(k'))^2}]. \tag{5.86}$$

$$\frac{\partial M_t^f(k)}{\partial L(i,j)} = M_{t-1}^f(i) Z_t^f(i,j) [\frac{\delta_{k',j}}{\sum_{k'} a_t^f(k')} - \frac{a_t^f(k)}{(\sum_{k'} a_t^f(k'))^2}]. \tag{5.87}$$

$$\begin{aligned} \frac{\partial W_t^f(i,j)}{\partial M_{t-1}^f(k)} &= (\delta_{i,k} L_t^f(i,j) Z(i,j) \\ &- W_t^f(i,j) [\sum_{j'} a_t^f(j') \frac{\partial M_t^f(j)}{\partial M_{t-1}^f(k)} + M_t^f(j) \sum_{j'} L_t^f(k,j') Z(k,j')]) / (M_t^f(j) \sum_{j'} a_t^f(j')). \end{aligned} \tag{5.88}$$

$$\begin{aligned} \frac{\partial W_t^f(i,j)}{\partial Z(k,l)} &= (\delta_{i,k} \delta_{j,l} L_t^f(i,j) M_{t-1}^f(i) \\ &- W_t^f(i,j) [\sum_{j'} a_t^f(j') \frac{\partial M_t^f(j)}{\partial Z(k,l)} + M_{t-1}^f(k) L_t^f(k,l) M_t^f(j)]) / (M_t^f(j) \sum_{j'} a_t^f(j')). \end{aligned} \tag{5.89}$$

$$\begin{aligned} \frac{\partial W_t^f(i,j)}{\partial L(k,l)} &= (\delta_{i,k} \delta_{j,l} Z(i,j) M_{t-1}^f(i) \\ &- W_t^f(i,j) [\sum_{j'} a_t^f(j') \frac{\partial M_t^f(j)}{\partial L(k,l)} + M_{t-1}^f(k) Z_t^f(k,l) M_t^f(j)]) / (M_t^f(j) \sum_{j'} a_t^f(j')). \end{aligned} \tag{5.90}$$

5.9.3 Analytical Derivatives of the Collapse Function

Recall the collapse operator: $(\mu_{t|t}(j), V_{t|t}(j)) = Collapse(\tilde{\mu}_{t|t}, \tilde{V}_{t|t}, W_t^f)$.

The corresponding derivatives are:

$$\frac{\partial \mu_{t|t}(i,j)}{\partial \tilde{\mu}_{t|t}[m]} = W(i,j)\delta_m \tag{5.91}$$

$$\frac{\partial \mu_{t|t}(i,j)}{\partial W_t[i]} = \tilde{\mu}_{t|t}(i,j) \tag{5.92}$$

$$\begin{aligned}\frac{\partial V_{t|t}(i,j)}{\partial \tilde{\mu}_{t|t}[m]} &= W(i,j)[(\delta_m - W(i,j)\delta_m)(\tilde{\mu}_{t|t}(i,j) - \mu_{t|t}(j))^\mathsf{T} \\ &+ (\tilde{\mu}_{t|t}(i,j) - \mu_{t|t}(j))(\delta_m - W(i,j)\delta_m)^\mathsf{T}] \\ &+ \sum_{k \neq i} W(k,j)[(-W(i,j)\delta_m)(\tilde{\mu}_{t|t}(k,j) - \mu_{t|t}(j))^\mathsf{T} \\ &+ (\tilde{\mu}_{t|t}(k,j) - \mu_{t|t}(j))(-W(i,j)\delta_m)^\mathsf{T}]\end{aligned} \tag{5.93}$$

$$\frac{\partial V_{t|t}(i,j)}{\partial \tilde{V}_{t|t}[m,n]} = \delta_{m,n}^{n,m} W(i,j) \tag{5.94}$$

$$\frac{\partial V_{t|t}(i,j)}{\partial W_t[i]} = \tilde{V}_{t|t}(i,j) + (\tilde{\mu}_{t|t}(i,j) - \mu_{t|t}(j))^\mathsf{T}(\tilde{\mu}_{t|t}(i,j) - \mu_{t|t}(j)) \tag{5.95}$$

5.10 Appendix B

5.10.1 Details of the Simulated Time-series

We simulated four different modes with the following parameters:

Mode 1:

$A_1 = [[[0.3,\ 0.5;\ -0.5,\ 0.5], [-0.4,\ 0.1;\ -0.2,\ 0.2]]; [eye(2,2),\ zeros(2,2)]]$

$C_1 = [1,\ 0,\ 0,\ 0;\ 0,\ 1,\ 0,\ 0]$

$Q_1 = [\sqrt{2},\ 0,\ 0,\ 0;\ 0, \sqrt{3},\ 0,\ 0;\ 0,\ 0,\ 1e-3,\ 0;\ 0,\ 0,\ 0,\ 1e-3]$

$R_1 = [0.05,\ 0;\ 0,\ 1]$

Mode 2:

$A_2 = [[[-0.3,\ -0.6;\ 0.4,\ -0.4], [0.3,\ -0.3;\ 0.4,\ -0.3]]; [zeros(2,2),\ zeros(2,2)]]$

$C_2 = [1,\ 0,\ 0,\ 0;\ 0,\ 1,\ 0,\ 0]$

$Q_2 = [\sqrt{2},\ 0,\ 0,\ 0;\ 0, \sqrt{3},\ 0,\ 0,\ ;\ 0,\ 0,\ 1e-3,\ 0;\ 0,\ 0,\ 0,\ 1e-3]$

$R_2 = [0.05,\ 0;\ 0,\ 1]$

Mode 3:

$A_3 = [[[-0.5,\ 0.2;\ 0.2,\ 0.5], [-0.4,\ 0.3;\ -0.4,\ 0.2]]; [eye(2,2),\ zeros(2,2)]]$

$C_3 = [1,\ 0,\ 0,\ 0;\ 0,\ 1,\ 0,\ 0,\ ]$

$Q_3 = [\sqrt{3},\ 0,\ 0,\ 0;\ 0\ , \sqrt{2},\ 0,\ 0;\ 0,\ 0,\ 1e-3,\ 0;\ 0,\ 0,\ 0,\ 1e-3]$

$R_3 = [1,\ 0;\ 0,\ 0.05]$

Mode 4:

$A_4 = [[[-0.4,\ 0.2;\ 0.2,\ -0.6], [0.5,\ -0.3;\ 0.4,\ -0.2]]; [eye(2,2),\ zeros(2,2)]]$

$C_4 = [1,\ 0,\ 0,\ 0;\ 0,\ 1,\ 0,\ 0]$

$Q_4 = [\sqrt{3},\ 0,\ 0,\ 0; 0, \sqrt{2},\ 0,\ 0\ ;\ 0,\ 0,\ 1e-3,\ 0;\ 0,\ 0,\ 0,\ 1e-3]$

$R_4 = [1,\ 0;\ 0,\ 0.05]$

The following four different mode transition matrices corresponding to four categories (A, B, C, and D) were defined (see Fig. 5-3):

Category A: Z_1 = [0.98, 0.005, 0.005, 0.01; 0.04, 0.92, 0.02, 0.02; 0.04, 0.02, 0.92, 0.02; 0.04, 0.02, 0.02, 0.92]

Category B: Z_2 = [0.92, 0.04, 0.02, 0.02; 0.01, 0.98, 0.005, 0.005; 0.01, 0.02, 0.96, 0.01; 0.04, 0.02, 0.02, 0.92]

Category C: Z_3 = [0.92, 0.02, 0.04, 0.02; 0.01, 0.96, 0.02, 0.01; 0.005, 0.005, 0.98, 0.01; 0.01, 0.01, 0.02, 0.96]

Category D: Z_4 = [0.92, 0.02, 0.02, 0.04; 0.02, 0.92, 0.02, 0.04; 0.01, 0.01, 0.96, 0.02; 0.005, 0.005, 0.01, 0.98]

All four models where initialized at $[0;0]$. Fig. 5-1 shows example time-series from each of the four categories. Using the parameters defined above, we simulated 200 time-series, with equal number per category. For each of the time-series, we simulated 600 samples and discarded the first 100 samples (to remove transients), resulting in 500 samples per time-series. Additionally, within each of the 200 simulated time-series, the upper block elements (corresponding to the AR coefficients) of the dynamics matrices $A_1, \cdots, A_4$ were perturbed by adding a white Gaussian noise with standard deviation of 0.05 to each of the elements. Therefore, although all the time-series shared the same modes, there were slight differences in dynamics.

Chapter 6

Conclusion and Future Work

Identification and analysis of physiological dynamical systems, in healthy and pathological states, may assist physiologists with better understanding of the etiology of complex diseases, and may aid clinicians with the development and titration of new treatments. As new patient monitoring and prognostic tools inspired by nonlinear dynamics are making their ways into the clinic (e.g., based on heart-rate entropy [117]), a systems physiology approach will require tools from multivariate nonlinear dynamical systems, where the interactions among physiological variables can be translated into a mechanistic description of the pathological state. The work presented in this thesis constitutes a step in bridging the gap between mathematical models of physiological control systems and the clinical utilization of such techniques, with the goal of improved patient state monitoring and clinical decision support.

6.1 Summary of Contributions

In **chapter 3** we investigated the feasibility of using fluctuations in multivariate physiological time-series to characterize (oscillatory) behaviors of the underlying systems around their resting stationary equilibrium points. Furthermore, we explored the utility of such characterization in quantifying system properties such as stability and the propensity to exhibit oscillatory outputs in response to external disturbances. We then developed an adaptive chemoreflex identification technique that incorporated measures of quality of experimentally recorded signals into the parameter estimation step, thus mitigating the influence of recording artifacts on the estimated model parameters. The potential clinical significance of this work includes the ability to assess respiratory instability in patients with sleep disordered breathing (e.g., Cheyne-Stokes breathing in congestive heart failure,

obstructive sleep apnea in adults, and neonatal apnea), and evaluation of weaning from mechanical ventilation in critically ill patients.

With the goal of modeling dynamical regime transitions in physiological time-series within a patient cohort, **Chapter 4** introduced the switching linear dynamical systems (SLDS) framework. We then presented an extension of the framework to incorporate physiological models of the underlying systems, and to discover similar dynamical patterns across a patient cohort. Our results from the tilt-table experiment demonstrated that the SLDS technique can correctly detect the switching dynamics in time-series of HR and BP, corresponding to various postural changes. Next, we applied the technique to a subset of patients from the MIMIC II intensive care unit database, where we showed that the evolving dynamics of time-series contain information pertaining to mortality of patients, both in-hospital as well as up to thirty days after hospital release.

In **Chapter 5** we developed an advanced machine-learning technique for outcome-discriminative learning of dynamics within a patient cohort. The main idea of our approach was to present the learning algorithm with the outcomes (or labels), corresponding to each time-series (e.g., survived vs. expired), and to learn time-series dynamics that are most relevant to the discriminative task of distinguishing among the labels. Using the proposed algorithm, we demonstrated a significant improvement in decoding postural changes involved in the tilt-table experiment, using the multivariate switching dynamics of HR and BP time-series. The technique developed in this chapter is also significant from a theoretical point of view, since we demonstrated that one may apply a backpropagation-based learning algorithm to the class of dynamic Bayesian networks if the goal of inference and learning is to use the learned marginal distributions for prediction or classification purposes.

6.2 Suggestions for Future Work

The techniques developed in this thesis allow for searching large physiological cohort time-series for dynamical patterns relevant to outcome variables of interest. Extending the work presented in this thesis, one may design efficient dynamics-based search algorithms for multivariate physiological time-series. Such tools will not only be useful for retrospective cohort studies (e.g., finding pathological dynamics in heart rate and blood pressure of patients at risk for hypotension) but also for online monitoring of patients and detection/prediction of significant clinical events (e.g., sepsis and onset of hypotension). A search tool customized for physiological time-series will have two

main components. First, offline algorithms for learning of multivariate dynamics, where the concept of "dynamics" is either defined by the user (e.g., based on constrained physiological models of the underlying systems) or is discovered by automated algorithms (i.e., system identification). This step will yield segmentation and indexing of time-series within the database. Next, given a bank of possible dynamical behaviors, the inference step will involve segmentation and assignment of newly presented patient time-series to the most likely dynamics, with the aim of event classification and predictive monitoring. Ultimately, tools from real-time monitoring of nonlinear dynamics in physiological time-series must be integrated with tools from the "-omics", bioimaging and other disciplines in order to develop comprehensive risk scores personalized to the individual patients.

Bibliography

[1] M. C. Mackey and J. G. Milton, "Dynamical diseases." *Ann N Y Acad Sci*, vol. 504, pp. 16–32, 1987.

[2] D. P. Francis, K. Willson, L. C. Davies, A. J. Coats, and M. Piepoli, "Quantitative general theory for periodic breathing in chronic heart failure and its clinical implications." *Circulation*, vol. 102, no. 18, pp. 2214–2221, Oct 2000.

[3] H. Rigatto and J. P. Brady, "Periodic breathing and apnea in preterm infants. ii. hypoxia as a primary event." *Pediatrics*, vol. 50, no. 2, pp. 219–228, Aug 1972.

[4] ——, "Periodic breathing and apnea in preterm infants. i. evidence for hypoventilation possibly due to central respiratory depression." *Pediatrics*, vol. 50, no. 2, pp. 202–218, Aug 1972.

[5] G. von Czettritz, R. T. Bax, T. Eckardt, S. Springer, and P. Emmrich, "[periodic breathing with periodic oxygen variation in infancy]," *Wien Med Wochenschr*, vol. 146, no. 13-14, pp. 317–319, 1996.

[6] M. H. Wilkinson, E. M. Skuza, G. C. Rennie, S. A. Sands, S. R. Yiallourou, R. S. C. Horne, and P. J. Berger, "Postnatal development of periodic breathing cycle duration in term and preterm infants." *Pediatr Res*, vol. 62, no. 3, pp. 331–336, Sep 2007. [Online]. Available: http://dx.doi.org/10.1203/PDR.0b013e3180db29e5

[7] S. Lahiri, K. Maret, and M. G. Sherpa, "Dependence of high altitude sleep apnea on ventilatory sensitivity to hypoxia." *Respir Physiol*, vol. 52, no. 3, pp. 281–301, Jun 1983.

[8] A. Berssenbrugge, J. Dempsey, C. Iber, J. Skatrud, and P. Wilson, "Mechanisms of hypoxia-induced periodic breathing during sleep in humans." *J Physiol*, vol. 343, pp. 507–524, Oct 1983.

[9] K. R. Burgess, P. L. Johnson, and N. Edwards, "Central and obstructive sleep apnoea during ascent to high altitude." *Respirology*, vol. 9, no. 2, pp. 222–229, Jun 2004. [Online]. Available: http://dx.doi.org/10.1111/j.1440-1843.2004.00576.x

[10] K. Burgess, K. Burgess, P. Subedi, P. Ainslie, Z. Topor, and W. Whitelaw, "Prediction of periodic breathing at altitude." *Adv Exp Med Biol*, vol. 605, pp. 442–446, 2008. [Online]. Available: http://dx.doi.org/10.1007/978-0-387-73693-8_77

[11] S. Javaheri, "A mechanism of central sleep apnea in patients with heart failure." *N Engl J Med*, vol. 341, no. 13, pp. 949–954, Sep 1999. [Online]. Available: http://dx.doi.org/10.1056/NEJM199909233411304

[12] T. D. Bradley and J. S. Floras, "Sleep apnea and heart failure: Part ii: central sleep apnea." *Circulation*, vol. 107, no. 13, pp. 1822–1826, Apr 2003. [Online]. Available: http://dx.doi.org/10.1161/01.CIR.0000061758.05044.64

[13] ——, "Sleep apnea and heart failure: Part i: obstructive sleep apnea." *Circulation*, vol. 107, no. 12, pp. 1671–1678, Apr 2003. [Online]. Available: http://dx.doi.org/10.1161/01.CIR.0000061757.12581.15

[14] A. Xie, J. B. Skatrud, D. S. Puleo, P. S. Rahko, and J. A. Dempsey, "Apnea-hypopnea threshold for co2 in patients with congestive heart failure." *Am J Respir Crit Care Med*, vol. 165, no. 9, pp. 1245–1250, May 2002.

[15] M. C. Khoo, R. E. Kronauer, K. P. Strohl, and A. S. Slutsky, "Factors inducing periodic breathing in humans: a general model." *J Appl Physiol*, vol. 53, no. 3, pp. 644–659, Sep 1982.

[16] T. Heldt, M. B. Oefinger, M. Hoshiyama, and R. G. Mark, "Circulatory response to passive and active changes in posture," *Computers in Cardiology*, vol. 30, p. 263266, 2003, circulatory response to passive and active changes in posture.

[17] T. Heldt, E. B. Shim, R. D. Kamm, and R. G. Mark, "Computational modeling of cardiovascular response to orthostatic stress." *J Appl Physiol*, vol. 92, no. 3, pp. 1239–1254, Mar 2002. [Online]. Available: http://dx.doi.org/10.1152/japplphysiol.00241.2001

[18] T. Parlikar and G. Verghese, "A simple cycle-averaged model for cardiovascular dynamics." *Conf Proc IEEE Eng Med Biol Soc*, vol. 5, pp. 5490–5494, 2005. [Online]. Available: http://dx.doi.org/10.1109/IEMBS.2005.1615726

[19] M. S. Olufsen, B. Smith, J. Mehlsen, and J. Ottesen, "The impact of gravity during head-up tilt." *Conf Proc IEEE Eng Med Biol Soc*, vol. 2011, pp. 2399–2402, 2011. [Online]. Available: http://dx.doi.org/10.1109/IEMBS.2011.6090669

[20] B. P. Grubb, "Clinical practice. neurocardiogenic syncope." *N Engl J Med*, vol. 352, no. 10, pp. 1004–1010, Mar 2005. [Online]. Available: http://dx.doi.org/10.1056/NEJMcp042601

[21] A. Porta, T. Bassani, V. Bari, E. Tobaldini, A. C. M. Takahashi, A. M. Catai, and N. Montano, "Model-based assessment of baroreflex and cardiopulmonary couplings during graded head-up tilt." *Comput Biol Med*, vol. 42, no. 3, pp. 298–305, Mar 2012. [Online]. Available: http://dx.doi.org/10.1016/j.compbiomed.2011.04.019

[22] P. K. Stein, P. P. Domitrovich, H. V. Huikuri, R. E. Kleiger, and C. I. , "Traditional and nonlinear heart rate variability are each independently associated with mortality after myocardial infarction." *J Cardiovasc Electrophysiol*, vol. 16, no. 1, pp. 13–20, Jan 2005.

[23] M. C. Khoo, F. Yang, J. J. Shin, and P. R. Westbrook, "Estimation of dynamic chemoresponsiveness in wakefulness and non-rapid-eye-movement sleep," *J Appl Physiol*, vol. 78, no. 3, pp. 1052–64, 1995, eng.

[24] R. Mukkamala, J. M. Mathias, T. J. Mullen, R. J. Cohen, and R. Freeman, "System identification of closed-loop cardiovascular control mechanisms: diabetic autonomic neuropathy," *Am J Physiol*, vol. 276, no. 3 Pt 2, pp. R905–12, 1999, eng.

[25] I. Korhonen, "Multivariate closed-loop model for analysis of cardiovascular dynamics," *Methods Inf Med*, vol. 36, no. 4-5, pp. 264–7, 1997, eng.

[26] A. Porta, T. Bassani, V. Bari, E. Tobaldini, A. C. Takahashi, A. M. Catai, and N. Montano, "Model-based assessment of baroreflex and cardiopulmonary couplings during graded head-up tilt," *Comput Biol Med*, vol. 42, no. 3, pp. 298–305, eng.

[27] P. C. Ivanov, M. G. Rosenblum, C. K. Peng, J. Mietus, S. Havlin, H. E. Stanley, and A. L. Goldberger, "Scaling behaviour of heartbeat intervals obtained by wavelet-based time-series analysis." *Nature*, vol. 383, no. 6598, pp. 323–327, Sep 1996. [Online]. Available: http://dx.doi.org/83323a0

[28] M. Costa, A. L. Goldberger, and C. K. Peng, "Multiscale entropy analysis of complex physiologic time series," *Phys Rev Lett*, vol. 89, no. 6, p. 068102, 2002, eng.

[29] M. Saeed, M. Villarroel, A. T. Reisner, G. Clifford, L.-W. Lehman, G. Moody, T. Heldt, T. H. Kyaw, B. Moody, and R. G. Mark, "Multiparameter intelligent monitoring in intensive care ii: a public-access intensive care unit database." *Crit Care Med*, vol. 39, no. 5, pp. 952–960, May 2011. [Online]. Available: http://dx.doi.org/10.1097/CCM.0b013e31820a92c6

[30] S. Kalli, J. Gronlund, H. Ihalainen, A. Siimes, I. Valimaki, and K. Antila, "Multivariate autoregressive modeling of autonomic cardiovascular control in neonatal lamb," *Comput Biomed Res*, vol. 21, no. 6, pp. 512–30, 1988, eng.

[31] R. Barbieri, G. Parati, and J. P. Saul, "Closed- versus open-loop assessment of heart rate baroreflex," *IEEE Engineering in Medicine and Biology Magazine*, vol. 20, no. 2, pp. 33–42, March–April 2001.

[32] S. D. Ghazanshahi and M. C. Khoo, "Estimation of chemoreflex loop gain using pseudorandom binary co2 stimulation." *IEEE Trans Biomed Eng*, vol. 44, no. 5, pp. 357–366, May 1997. [Online]. Available: http://dx.doi.org/10.1109/10.568911

[33] S. Nemati, B. A. Edwards, S. A. Sands, P. J. Berger, A. Wellman, G. C. Verghese, A. Malhotra, and J. P. Butler, "Model-based characterization of ventilatory stability using spontaneous breathing," *J Appl Physiol*, vol. 111, no. 1, pp. 55–67, 2011, eng.

[34] L. Ljung, *System identification toolbox: for use with MATLAB: user's guide: version 5*. MathWorks, Inc, 1997.

[35] A. Schlögl, "A comparison of multivariate autoregressive estimators," *Signal Processing*, vol. 86, no. 9, pp. 2426 – 2429, 2006, ¡ce:title¿Special Section: Signal Processing in UWB Communications¡/ce:title¿. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0165168405003993

[36] S. Mason, "Feedback theory: further properties of signal flow graphs," *Technical report (Massachusetts Institute of Technology. Research Laboratory of Electronics)*, 1956.

[37] M. Priestley, *Spectral analysis and time series (Probability and mathematical statistics)*. Academic Press Limited, London, 1994.

[38] J. Lardes, "Relationship between state-space and armav approaches to modal parameter identification," *Mechanical Systems and Signal Processing*, vol. 22, no. 3, pp. 611–616, 2008, relationship between state-space and ARMAV approaches to modal parameter identification.

[39] Z. Ghahramani, S. Roweis, *et al.*, "Learning nonlinear dynamical systems using an em algorithm," *Advances in neural information processing systems*, pp. 431–437, 1999.

[40] R. Brown, P. Hwang, *et al.*, *Introduction to random signals and applied Kalman filtering*. Wiley New York, 1992, vol. 1.

[41] V. Digalakis, J. Rohlicek, and M. Ostendorf, "Ml estimation of a stochastic linear system with the em algorithm and its application to speech recognition," *Speech and Audio Processing, IEEE Transactions on*, vol. 1, no. 4, pp. 431–442, 1993.

[42] K. P. Murphy, "Switching kalman filter," *Compaq Cambridge Research Laboratory, Tech. Rep. 98-10.*, 1998, cambridge, MA.

[43] T. Katayama, *Subspace methods for system identification*. Springer, 2005.

[44] B. Boots, "Learning stable linear dynamical systems."

[45] Z. Ghahramani and G. Hinton, "Parameter estimation for linear dynamical systems," Technical Report CRG-TR-96-2, University of Totronto, Dept. of Computer Science, Tech. Rep., 1996.

[46] G. C. Verghese, I. J. Perez-Arriaga, and F. C. Schweppe, "Selective modal analysis with applications to electric power systems, part ii: The dynamic stability problem," *Power Apparatus and Systems, IEEE Transactions on*, vol. PAS-101, no. 9, pp. 3126–3134, 1982, 0018-9510.

[47] I. Perez-Arriaga and F. Verghese, G. & Schweppe, "Selective modal analysis with applications to electric power systems, part i: Heuristic introduction," *EEE Transactions on Power Apparatus and Systems*, no. 9, pp. 3117 – 3125, 1982.

[48] M. Arnold, W. H. Miltner, H. Witte, R. Bauer, and C. Braun, "Adaptive ar modeling of nonstationary time series by means of kalman filtering." *IEEE Trans Biomed Eng*, vol. 45, no. 5, pp. 553–562, May 1998. [Online]. Available: http://dx.doi.org/10.1109/10.668741

[49] M. J. Cassidy and W. D. Penny, "Bayesian nonstationary autoregressive models for biomedical signal analysis." *IEEE Trans Biomed Eng*, vol. 49, no. 10, pp. 1142–1152, Oct 2002. [Online]. Available: http://dx.doi.org/10.1109/TBME.2002.803511

[50] M. P. Tarvainen, S. D. Georgiadis, P. O. Ranta-Aho, and P. A. Karjalainen, "Time-varying analysis of heart rate variability signals with a kalman smoother algorithm." *Physiol Meas*, vol. 27, no. 3, pp. 225–239, Mar 2006. [Online]. Available: http://dx.doi.org/10.1088/0967-3334/27/3/002

[51] S. Nemati, A. Malhotra, and G. D. Clifford, "Data fusion for improved respiration rate estimation." *EURASIP J Adv Signal Process*, vol. 2010, p. 926305, 2010. [Online]. Available: http://dx.doi.org/10.1155/2010/926305

[52] B. Edwards, S. Sands, E. Skuza, E. Stockx, V. Brodecky, M. Wilkinson, and P. Berger, "Increased peripheral chemosensitivity via dopaminergic manipulation promotes respiratory instability in lambs," *Respiratory physiology & neurobiology*, vol. 164, no. 3, pp. 419–428, 2008.

[53] M. C. Khoo and V. Z. Marmarelis, "Estimation of peripheral chemoreflex gain from spontaneous sigh responses." *Ann Biomed Eng*, vol. 17, no. 6, pp. 557–570, 1989.

[54] M. Younes, M. Ostrowski, R. Atkar, J. Laprairie, A. Siemens, and P. Hanly, "Mechanisms of breathing instability in patients with obstructive sleep apnea." *J Appl Physiol*, vol. 103, no. 6, pp. 1929–1941, Dec 2007. [Online]. Available: http://dx.doi.org/10.1152/japplphysiol.00561.2007

[55] A. Wellman, A. Malhotra, A. S. Jordan, K. E. Stevenson, S. Gautam, and D. P. White, "Effect of oxygen in obstructive sleep apnea: role of loop gain." *Respir Physiol Neurobiol*, vol. 162, no. 2, pp. 144–151, Jul 2008. [Online]. Available: http://dx.doi.org/10.1016/j.resp.2008.05.019

[56] I. S. Breslav and A. D. Nozdrachev, "[regulation of respiration: visceral and behavioural components]," *Usp Fiziol Nauk*, vol. 38, no. 2, pp. 26–45, 2007.

[57] M. J. Mador and M. J. Tobin, "Effect of alterations in mental activity on the breathing pattern in healthy subjects." *Am Rev Respir Dis*, vol. 144, no. 3 Pt 1, pp. 481–487, Sep 1991.

[58] E. A. Phillipson, "Control of breathing during sleep." *Am Rev Respir Dis*, vol. 118, no. 5, pp. 909–939, Nov 1978.

[59] N. S. Cherniack and G. S. Longobardo, "Cheyne-stokes breathing. an instability in physiologic control." *N Engl J Med*, vol. 288, no. 18, pp. 952–957, May 1973.

[60] M. Younes, M. Ostrowski, W. Thompson, C. Leslie, and W. Shewchuk, "Chemical control stability in patients with obstructive sleep apnea." *Am J Respir Crit Care Med*, vol. 163, pp. 1181–1190, Apr 2001.

[61] J. A. Dempsey, C. A. Smith, T. Przybylowski, B. Chenuel, A. Xie, H. Nakayama, and J. B. Skatrud, "The ventilatory responsiveness to co(2) below eupnoea as a determinant of ventilatory stability in sleep." *J Physiol*, vol. 560, no. Pt 1, pp. 1–11, Oct 2004. [Online]. Available: http://dx.doi.org/10.1113/jphysiol.2004.072371

[62] L. Goodman, "Oscillatory behavior of ventilation in resting man." *IEEE Trans Biomed Eng*, vol. 11, pp. 82–93, Jul 1964.

[63] E. N. Bruce, "Temporal variations in the pattern of breathing." *J Appl Physiol*, vol. 80, no. 4, pp. 1079–1087, Apr 1996.

[64] M. Modarreszadeh and E. N. Bruce, "Ventilatory variability induced by spontaneous variations of paco2 in humans." *J Appl Physiol*, vol. 76, no. 6, pp. 2765–2775, Jun 1994.

[65] J. G. V. den Aardweg and J. M. Karemaker, "Influence of chemoreflexes on respiratory variability in healthy subjects." *Am J Respir Crit Care Med*, vol. 165, no. 8, pp. 1041–1047, Apr 2002.

[66] N. S. Cherniack, G. S. Longobardo, O. R. Levine, R. Mellins, and A. P. Fishman, "Periodic breathing in dogs." *J Appl Physiol*, vol. 21, no. 6, pp. 1847–1854, Nov 1966.

[67] M. H. Wilkinson, P. J. Berger, N. Blanch, V. Brodecky, and C. Jones, "Source of respiratory drive during periodic breathing in lambs." *Respir Physiol*, vol. 104, no. 2-3, pp. 115–126, Jul 1996.

[68] M. H. Wilkinson, K.-L. Sia, E. M. Skuza, V. Brodecky, and P. J. Berger, "Impact of changes in inspired oxygen and carbon dioxide on respiratory instability in the lamb." *J Appl Physiol*, vol. 98, no. 2, pp. 437–446, Feb 2005. [Online]. Available: http://dx.doi.org/10.1152/japplphysiol.00532.2004

[69] R. Oguma, "Method to evaluate signal transmission paths in dynamic systems and its application to reactor noise analysis," *Journal of Nuclear Science and Technology*, vol. 18, no. 11, pp. 835–844, 1981.

[70] B. A. Edwards, S. A. Sands, E. M. Skuza, V. Brodecky, E. M. Stockx, M. H. Wilkinson, and P. J. Berger, "Maturation of respiratory control and the propensity for breathing instability in a sheep model." *J Appl Physiol*, vol. 107, no. 5, pp. 1463–1471, Nov 2009. [Online]. Available: http://dx.doi.org/10.1152/japplphysiol.00587.2009

[71] G. D. Mitsis, R. J. M. Governo, R. Rogers, and K. T. S. Pattinson, "The effect of remifentanil on respiratory variability, evaluated with dynamic modeling," *J Appl Physiol*, vol. 106, no. 4, pp. 1038–1049, Apr 2009. [Online]. Available: http://dx.doi.org/10.1152/japplphysiol.90769.2008

[72] A. Jubran, B. J. Grant, and M. J. Tobin, "Effect of hyperoxic hypercapnia on variational activity of breathing." *Am J Respir Crit Care Med*, vol. 156, no. 4 Pt 1, pp. 1129–1139, Oct 1997. [Online]. Available: http://ajrccm.atsjournals.org/cgi/reprint/156/4/1129

[73] B. A. Edwards, S. A. Sands, E. M. Skuza, E. M. Stockx, V. Brodecky, M. H. Wilkinson, and P. J. Berger, "Increased peripheral chemosensitivity via dopaminergic manipulation promotes respiratory instability in lambs," *Respir Physiol Neurobiol*, vol. 164, no. 3, pp. 419–28, 2008, eng.

[74] A. Chilton and R. Stacy, "A mathematical analysis of carbon dioxide respiration in man," *Bulletin of Mathematical Biology*, vol. 14, no. 1, pp. 1–18, 1952.

[75] M. H. Wilkinson, P. J. Berger, N. Blanch, and V. Brodecky, "Effect of venous oxygenation on arterial desaturation rate during repetitive apneas in lambs." *Respir Physiol*, vol. 101, no. 3, pp. 321–331, Sep 1995.

[76] J. P. Miller, D. J. Cunningham, B. B. Lloyd, and J. M. Young, "The transient respiratory effects in man of sudden changes in alveolar co2 in hypoxia and in high oxygen." *Respir Physiol*, vol. 20, no. 1, pp. 17–31, Feb 1974.

[77] A. Berkenbosch, A. Dahan, J. DeGoede, and I. C. Olievier, "The ventilatory response to co2 of the peripheral and central chemoreflex loop before and after sustained hypoxia in man." *J Physiol*, vol. 456, pp. 71–83, Oct 1992.

[78] M. E. Pedersen, M. Fatemian, and P. A. Robbins, "Identification of fast and slow ventilatory responses to carbon dioxide under hypoxic and hyperoxic conditions in humans." *J Physiol*, vol. 521 Pt 1, pp. 273–287, Nov 1999. [Online]. Available: http://jp.physoc.org/cgi/content/abstract/521/1/273?ck=nck

[79] C. A. Smith, J. R. Rodman, B. J. A. Chenuel, K. S. Henderson, and J. A. Dempsey, "Response time and sensitivity of the ventilatory response to co2 in unanesthetized intact dogs: central vs. peripheral chemoreceptors." *J Appl Physiol*, vol. 100, no. 1, pp. 13–19, Jan 2006. [Online]. Available: http://dx.doi.org/10.1152/japplphysiol.00926.2005

[80] R. A. Henry, I. L. Lu, L. A. Beightol, and D. L. Eckberg, "Interactions between co2 chemoreflexes and arterial baroreflexes." *Am J Physiol*, vol. 274, no. 6 Pt 2, pp. H2177–H2187, Jun 1998.

[81] P. Solin, T. Roebuck, D. P. Johns, E. H. Walters, and M. T. Naughton, "Peripheral and central ventilatory responses in central sleep apnea with and without congestive heart failure." *Am J Respir Crit Care Med*, vol. 162, no. 6, pp. 2194–2200, Dec 2000.

[82] P. Ponikowski, T. P. Chua, S. D. Anker, D. P. Francis, W. Doehner, W. Banasiak, P. A. Poole-Wilson, M. F. Piepoli, and A. J. Coats, "Peripheral chemoreceptor hypersensitivity: an ominous sign in patients with chronic heart failure." *Circulation*, vol. 104, no. 5, pp. 544–549, Jul 2001.

[83] P. A. Lanfranchi, A. Braghiroli, E. Bosimini, G. Mazzuero, R. Colombo, C. F. Donner, and P. Giannuzzi, "Prognostic value of nocturnal cheyne-stokes respiration in chronic heart failure." *Circulation*, vol. 99, no. 11, pp. 1435–1440, Mar 1999.

[84] E. N. Bruce, "Assessing respiratory control during spontaneous breathing: practice may be more difficult than theory." *Am J Respir Crit Care Med*, vol. 165, no. 8, pp. 1033–1034, Apr 2002.

[85] M. C. Khoo, ""ventilatory stability to co2 disturbances in wakefulness and quiet sleep"." *J Appl Physiol*, vol. 79, no. 4, pp. 1069–1070, Oct 1995.

[86] M. Modarreszadeh, E. N. Bruce, H. Hamilton, and D. W. Hudgel, "Ventilatory stability to co2 disturbances in wakefulness and quiet sleep." *J Appl Physiol*, vol. 79, no. 4, pp. 1071–1081, Oct 1995.

[87] A. Loewen, M. Ostrowski, J. Laprairie, R. Atkar, J. Gnitecki, P. Hanly, and M. Younes, "Determinants of ventilatory instability in obstructive sleep apnea: inherent or acquired?" *Sleep*, vol. 32, no. 10, pp. 1355–1365, Oct 2009.

[88] A. Salloum, J. A. Rowley, J. H. Mateika, S. Chowdhuri, Q. Omran, and M. S. Badr, "Increased propensity for central apnea in patients with obstructive sleep apnea: effect of nasal continuous positive airway pressure." *Am J Respir Crit Care Med*, vol. 181, no. 2, pp. 189–193, Jan 2010. [Online]. Available: http://dx.doi.org/10.1164/rccm.200810-1658OC

[89] M. C. Khoo, "Determinants of ventilatory instability and variability." *Respir Physiol*, vol. 122, no. 2-3, pp. 167–182, Sep 2000.

[90] H. Akaike, "A new look at the statistical model identification," *IEEE transactions on automatic control*, vol. 19, no. 6, pp. 716–723, 1974.

[91] G. Benchetrit and F. Bertrand, "A short-term memory in the respiratory centres: statistical analysis." *Respir Physiol*, vol. 23, no. 2, pp. 147–158, Mar 1975.

[92] P. J. Liang, J. J. Pandit, and P. A. Robbins, "Statistical properties of breath-to-breath variations in ventilation at constant petco2 and peto2 in humans." *J Appl Physiol*, vol. 81, no. 5, pp. 2274–2286, Nov 1996.

[93] I. D. Clement and P. A. Robbins, "Latency of the ventilatory chemoreflex response to hypoxia in humans." *Respir Physiol*, vol. 92, no. 3, pp. 277–287, Jun 1993.

[94] M. Ding, S. Bressler, W. Yang, and H. Liang, "Short-window spectral analysis of cortical event-related potentials by adaptive multivariate autoregressive modeling: data preprocessing, model validation, and variability assessment," *Biological cybernetics*, vol. 83, no. 1, pp. 35–45, 2000.

[95] S. Nemati, L. H. Lehman, R. P. Adams, and M. A., "Discovering shared cardiovascular dynamics within a patient cohort," *Engineering in Medicine and Biology Society (EMBC)*, 2012, accepted.

[96] G. Parati, J. P. Saul, M. Di Rienzo, and G. Mancia, "Spectral analysis of blood pressure and heart rate variability in evaluating cardiovascular regulation. a critical appraisal," *Hypertension*, vol. 25, no. 6, pp. 1276–86, 1995, eng.

[97] M. Di Rienzo, G. Parati, G. Mancia, A. Pedotti, and P. Castiglioni, "Investigating baroreflex control of circulation using signal processing techniques," *Engineering in Medicine and Biology Magazine, IEEE*, vol. 16, no. 5, pp. 86–95, 1997, investigating baroreflex control of circulation using signal processing techniques.

[98] P. C. Ivanov, L. A. Amaral, A. L. Goldberger, S. Havlin, M. G. Rosenblum, Z. R. Struzik, and H. E. Stanley, "Multifractality in human heartbeat dynamics," *Nature*, vol. 399, no. 6735, pp. 461–5, 1999, eng.

[99] M. Saeed and R. Mark, "A novel method for the efficient retrieval of similar multiparameter physiologic time series using wavelet-based symbolic representations," *AMIA Annu Symp Proc*, pp. 679–83, 2006, eng.

[100] R. K. Avent and J. D. Charlton, "A critical review of trend-detection methodologies for biomedical monitoring systems," *Crit Rev Biomed Eng*, vol. 17, no. 6, pp. 621–59, 1990, eng.

[101] J. R. Le Gall, P. Loirat, A. Alperovitch, P. Glaser, C. Granthil, D. Mathieu, P. Mercier, R. Thomas, and D. Villers, "A simplified acute physiology score for icu patients." *Crit Care Med*, vol. 12, no. 11, pp. 975–977, Nov 1984.

[102] J. R. Le Gall, S. Lemeshow, and F. Saulnier, "A new simplified acute physiology score (saps ii) based on a european/north american multicenter study." *JAMA*, vol. 270, no. 24, pp. 2957–2963, 1993.

[103] J. E. Zimmerman, A. A. Kramer, D. S. McNair, and F. M. Malila, "Acute physiology and chronic health evaluation (apache) iv: hospital mortality assessment for today's critically ill patients." *Crit Care Med*, vol. 34, no. 5, pp. 1297–1310, May 2006. [Online]. Available: http://dx.doi.org/10.1097/01.CCM.0000215112.84523.F0

[104] A. L. Goldberger, L. A. Amaral, J. M. Hausdorff, P. Ivanov, C. K. Peng, and H. E. Stanley, "Fractal dynamics in physiology: alterations with disease and aging," *Proc Natl Acad Sci U S A*, vol. 99 Suppl 1, pp. 2466–72, 2002, eng.

[105] A. C. Fowler and M. J. McGuinness, "A delay recruitment model of the cardiovascular control system," *J Math Biol*, vol. 51, no. 5, pp. 508–26, 2005, eng.

[106] P. McSharry, M. McGuinness, and A. Fowler, "Confronting a cardiovascular system model with heart rate and blood pressure data," *Computers in Cardiology*, pp. 587–590, 2005, confronting a cardiovascular system model with heart rate and blood pressure data.

[107] "Heart rate variability: standards of measurement, physiological interpretation and clinical use. task force of the european society of cardiology and the north american society of pacing and electrophysiology," *Circulation*, vol. 93, no. 5, pp. 1043–65, 1996, eng.

[108] L. A. Celi, L. C. Hinske, G. Alterovitz, and P. Szolovits, "An artificial intelligence tool to predict fluid requirement in the intensive care unit: a proof-of-concept study," *Crit Care*, vol. 12, no. 6, p. R151, 2008, eng.

[109] E. Fox, E. Sudderth, M. Jordan, and A. Willsky, "Sharing features among dynamical systems with beta processes," *Advances in Neural Information Processing Systems*, vol. 22, pp. 549–557, 2009.

[110] L. Lehman, S. Nemati, R. Adams, and R. Mark, "Discovering shared dynamics in physiological signals: Application to patient monitoring in icu," *Engineering in Medicine and Biology Society (EMBC)*, 2012, accepted. [Online]. Available: http://hips.seas.harvard.edu/files/lehman-icu-embc-2012.pdf

[111] G. Hinton, S. Osindero, and Y. Teh, "A fast learning algorithm for deep belief nets," *Neural computation*, vol. 18, no. 7, pp. 1527–1554, 2006.

[112] D. Erhan, Y. Bengio, A. Courville, P. Manzagol, P. Vincent, and S. Bengio, "Why does unsupervised pre-training help deep learning?" *The Journal of Machine Learning Research*, vol. 11, pp. 625–660, 2010.

[113] R. Memisevic, "An introduction to structured discriminative learning," Citeseer, Tech. Rep., 2006.

[114] F. Eaton and Z. Ghahramani, "Choosing a variable to clamp: approximate inference using conditioned belief propagation," *n Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics I*, vol. 5, pp. 145–152, 2009.

[115] V. Stoyanov, A. Ropson, and J. Eisner, "Empirical risk minimization of graphical model parameters given approximate inference, decoding, and model structure," in *Proceedings of AISTATS*, 2011.

[116] J. Domke, "Beating the likelihood: Marginalization-based parameter learning in graphical models."

[117] J. R. Moorman, J. B. Delos, A. A. Flower, H. Cao, B. P. Kovatchev, J. S. Richman, and D. E. Lake, "Cardiovascular oscillations at the bedside: early diagnosis of neonatal sepsis using heart rate characteristics monitoring." *Physiol Meas*, vol. 32, no. 11, pp. 1821–1832, Nov 2011. [Online]. Available: http://dx.doi.org/10.1088/0967-3334/32/11/S08