State-space Decoders for Wearable Healthcare Applications

by
Dilranjan S. Wickramasuriya

A dissertation submitted to the Department of Electrical and Computer Engineering,

College of Engineering

in partial fulfillment of the requirements for the degree of

Doctor of Philosophy
in Electrical Engineering

Chair of Committee: Rose T. Faghih

Committee Member: Jose L. Contreras-Vidal

Committee Member: Zhu Han

Committee Member: Joseph T. Francis

Committee Member: Hien V. Nguyen

University of Houston
December 2020

DEDICATION/EPIGRAPH

To Chris – a support, an example and most importantly, my friend

ACKNOWLEDGMENTS

I would first like to express my sincere gratitude to my advisor Dr. Rose Faghih for her guidance and supervision throughout the course of the past three and a half years. I am also grateful to members on my PhD committee for helpful advice and suggestions. I also wish to acknowledge the financial support of the National Science Foundation (NSF), in particular, my study was supported by the following NSF grants: 1755780 – CRII: CPS: Wearable-Machine Interface Architectures and 1942585 – CAREER: MINDWATCH: Multimodal Intelligent Noninvasive brain state Decoder for Wearable AdapTive Closed-loop arcHitectures. Finally, I am thankful to friends and family who have been with me on this journey at the University of Houston.

ABSTRACT

Homeostatic processes govern multiple latent state variables within the human body. While many of these states remain largely unobserved, they frequently give rise to bioelectric and biochemical phenomena that can be measured. The measured signals provide a window into estimating the unobserved states. In a number of instances, the observed electrical and chemical phenomena are pulsatile or impulse-like in nature. This dissertation describes state-space methods for estimating latent variables tied to changes in skin conductance, heart rate and cortisol secretion – all of which have a characteristic point process nature.

In the first two sections, state-space methods are developed for estimating sympathetic arousal from skin conductance and heart rate features. Estimation involves Bayesian filtering applied within an expectation-maximization framework. Results are provided on experiments involving different types of mental stressors and Pavlovian fear conditioning. The results agree with general expectations with high arousal levels typically occurring during stressors and lower values occurring during relaxation. General agreement with expectations is also found with different trial averages in fear conditioning. Skin conductance-based estimates are also validated with blood flow signals in the brain in a separate experiment.

In the third section, state-space methods are developed to estimate energy production from blood cortisol measurements. The methods are applied to simulated and experimental data from patients suffering from Cushing's disease, chronic fatigue syndrome and fibromyalgia syndrome. The results help shed light on why patients with hypercortisolism may experience daytime fatigue and nighttime sleeping difficulties. Circadian-like behavior is also seen with higher energy estimates occurring towards morning awakening and lower values at bedtime.

In the final section, machine learning methods are used for state-space estimation. Traditional Bayesian filtering methods do not have the ability to permit external influences such as domain knowledge or labels to affect the state estimates. We develop a hybrid estimator that enables this possibility and apply it to both skin conductance-based arousal estimation and cortisol-related energy estimation. The hybrid estimator permits the enforcement of

circadian rhythms on to the state estimates and the customization of the level to which the external influence is permitted to affect them.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1 Introduction

Annual healthcare expenditure in the United States exceeded $3.5 trillion in 2017 [1], and is expected to grow even larger over the coming decade [2]. Moreover, healthcare spending *per capita* in the United States is one of the highest in the developed world [3]. While many solutions have been proposed to curb ballooning healthcare expenses—including changes to government policy, insurance structure and drug pricing—concrete steps to stem the upward spiral of costs are yet to materialize. A possible technological breakthrough to impede rising healthcare costs lies in the development of intelligent healthcare systems that can care for patients. Here, the point-of-care would shift back to the patient. Currently, the delivery of healthcare is largely physician- or hospital-centric. Individuals typically wait until they fall ill and then proceed to obtain treatment in a clinical setting. With smart healthcare systems in place, the emphasis will be on utilizing wearable monitoring, signal analysis tools and domain knowledge to help patients manage their own clinical needs, or even just help keep people well. This gradual transition to *mHealth* is one that holds much promise [4, 5].

Wearable monitoring is thus expected to play a critical role in the future of healthcare [6, 7]. Consider the representative graphic of a wearable healthcare system in Fig. 1. Different sensors can be used to monitor the patient. The sensors may be in the form of smartwatches, chemical sensor patches that analyze body fluids or smart textiles with integrated sensing capabilities. The physiological signals recorded from the patient are primarily electrical or chemical. The bioelectric signals may include electroencephalography (EEG), skin conductance, electrocardiography (EKG), photoplethysmography (PPG) or electromyography (EMG). Hormone concentrations in the blood, salt levels in the sweat and biochemical saliva compositions are examples of the chemical signals that may be monitored. The signals acquired from these sensors tell us something about the *underlying states* within the body. These states could be related to emotion, metabolic energy production or disease.

The states, however, cannot be observed directly. Nevertheless, biomarkers in the signals that are continuously monitored provide information regarding them. The signals can thus be analyzed and corrective action prescribed based on the current state or condition at hand. The corrective action may involve prescribing a short-acting drug, adjusting the parameters of an implanted neurostimulation device or sending a simple text alert suggesting a walk outdoors. This, we argue, is a closed-loop control system—one that seeks to regulate a particular aspect of health or well-being. Now the problem of sensing latent states in a system, estimating them and finally applying corrective control is one that has been investigated for decades in controls engineering. As such, tools from controls theory can be readily adapted to the case of designing the intelligent closed-loop man-machine healthcare systems of the future.

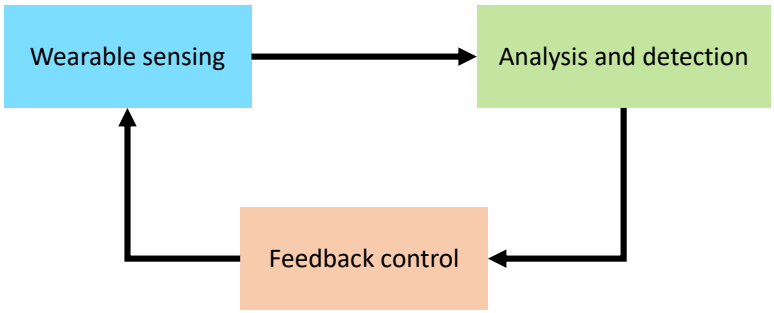

Figure 1: Closed-loop medical therapy.

Here we develop decoders based on state-space models for wearable healthcare applications in two different scenarios. The first concerns the estimation of sympathetic arousal from skin conductance and heart rate measurements, and the second concerns the estimation of energy production from blood cortisol measurements. Currently, automated closed-loop neurostimulation therapies for treating psychiatric disorders and automated cortisol infusion pumps are lacking. The state-space formulations presented herein could be a step towards their eventual future realization.

1.1 Disorders of Psychological Arousal

According to the National Alliance on Mental Health, 19.1% of adults in the United States (approximately 47.6 million) experienced mental illness in 2018 [8]. The illnesses were serious in about a quarter of the cases. The annual prevalence of a major depressive episode was 7.2% and that of post-traumatic stress disorder (PTSD) was 3.9% that year. Depression and PTSD are, in some ways, on opposite ends of a spectrum. Depression typically involves *low* levels of psychological arousal [9]. James A. Russell, in his influential 1980 paper entitled "A Circumplex Model of Affect," described a model where emotion could be accounted for along two orthogonal axes—valence and arousal [10]. Valence denotes the pleasant-unpleasant nature of an emotion while arousal captures its corresponding activation or excitement. In contrast to depression, PTSD patients often experience symptoms of *hyper*arousal that include "being easily startled, feeling tense, having difficulty sleeping and/or having angry outbursts" [11]. While PTSD and depression are associated with symptoms on opposite ends of the arousal spectrum, bipolar disorder (also known as manic depression) is characterized by unusual mood swings that range between the two, i.e., feelings vary from extreme emotional highs to depression [12]. Note also that the term "emotional" and "affective" are used interchangeably in the literature [13]. For instance, the valence-arousal model in [10] is called the Circumplex Model of *Affect*. Arousal is tied to the level of activation of the sympathetic branch of the autonomic nervous system. Consequently, phrases such as "emotional arousal," "sympathetic arousal", "autonomic arousal" and "affective arousal" are all used in the literature. Technically speaking, however, sympathetic arousal is a broader term governing the body's fight-or-flight response and emotional arousal is one part of it [14, 15].

While depression and PTSD can be treated with medication and/or psychotherapy, the symptoms do not diminish with treatment in certain cases. Neurostimulation therapies such as deep brain stimulation (DBS) are then a viable alternative, both for depression [16, 17] and for PTSD [18, 19]. Most commercial DBS systems function in a open-loop manner.

A patient visits a physician periodically who manually adjusts electrical stimulation parameters based on current symptoms. Closed-loop DBS (CLDBS) systems are emerging as the next generation of the technology where stimulation will adjust automatically instead. CLDBS has been shown to have superior performance in several studies investigating Parkinson's disease [20, 21, 22]. CLDBS for treating disorders of psychological arousal currently do not exist. Due to its sensitive nature to psychological arousal [23], skin conductance can be used to estimate arousal to help close the loop in adaptive DBS systems for treating PTSD or depression [24] (Fig. 2). The applications of skin conductance-based arousal detection are, however, not just limited to PTSD or depression alone, nor to CLDBS systems. Patients suffering from other disorders of psychological arousal could also benefit from it (e.g., in the case of a patient suffering from bipolar disorder, both emotional highs and lows could be detected to monitor the patient's progress or recovery over time). The approach could also be used for general stress management and maintaining emotional well-being.

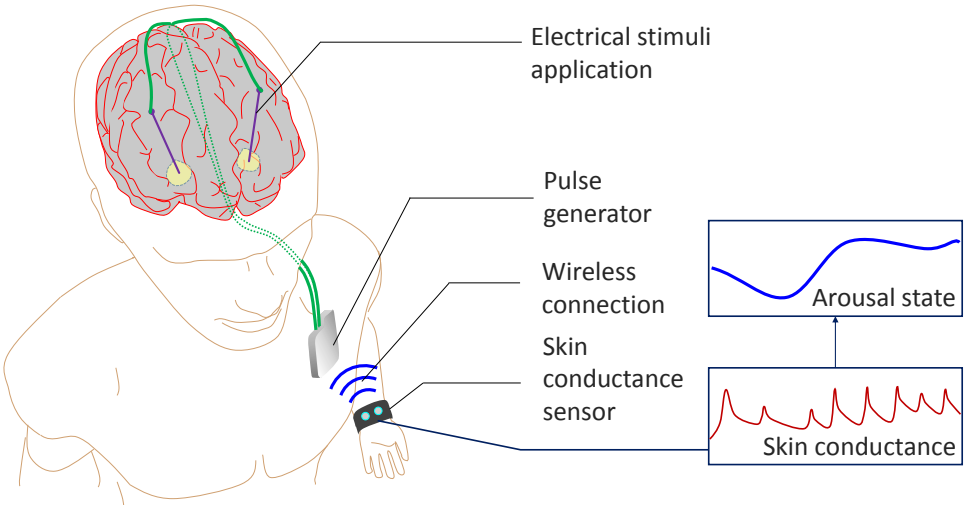

Figure 2: Adaptive closed-loop neurostimulation based on measuring skin conductance.

1.2 Cortisol Dysregulation

Cortisol is the body's main stress hormone. One of its primary purposes is to raise blood glucose levels in response to external stressors [25, 26]. As such, it is categorized among the class of hormones known as glucocorticoids. Cortsiol secretion is governed by a feedback control system involving the hypothalamus, pituitary and adrenal glands. These, together with feedback signals, messaging mechanisms and the resulting hormone secretions form what is commonly known as the hypothalamic-pituitary-adrenal (HPA) axis. Cortisol has other functions including the suppression of the immune response and playing a role in certain metabolic processes. Two types of disorders are typically associated with cortisol—hypocortisolism (too little cortisol in the bloodstream) and hypercortisolism (too much cortisol). Now cortisol is secreted in pulses [27, 28]. Typically, between 15–22 pulses of cortisol are secreted in a healthy adult every day [29, 30]. Cortisol disorders are mostly treated with oral medication. For instance, in Addison's disease—a type of hypocortisolism involving adrenal gland dysfunction—patients are treated with one or two doses of hydrocortisone daily. This course of treatment, however, may not be optimal [31]. Concerning the treatment of congenital adrenal hyperplasia (another type of adrenal deficiency), Nella et al., [32] point out that, "conventional cortisol replacement with oral glucocorticoids once, twice, or thrice daily remains suboptimal." Given this lack of optimality, an automated closed-loop infusion system for treating cortisol disorders in a manner that mimics the body's own secretory mechanism would be more desirable. Such infusion pumps are already a topic of investigation in the literature.

Automated closed-loop *insulin* pumps for treating patients with Diabetes are now marketed by Medtronic (e.g., the MiniMed 670G Insulin Pump System). Devices in this family continuously monitor blood glucose levels and automatically infuse insulin based on a control algorithm. However, the corresponding *cortisol* infusion pump for treating cortisol dysregulation currently does not exist. Instead, preliminary studies investigating continuous automated cortisol infusion have relied on an off-label use of a device in the class of

Medtronic's insulin pumps (e.g., [32, 33, 34]). Encouraging results from these investigations have shown the superiority of continuous infusion over intermittent doses of oral medication. Unfortunately, a commercial sensor for convenient long-term cortisol monitoring is currently unavailable. Without a such a sensor, off-label use of the Medtronic insulin pumps for cortisol infusion utilize algorithms that mimic the body's circadian variation of cortisol. As pointed out before, cortisol is released in pulses. A system-theoretic formulation for the infusion of pulsatile cortisol into the bloodstream based on sensor measurements that could be embedded within an automated pump is presently lacking.

Thus we envisage two primary closed-loop application scenarios for the methods that will be described in this dissertation. The first relates to the treatment of psychological arousal disorders using automated neurostimulation therapies and second to the automated infusion of cortisol (hydrocortisone) for treating cortisol disorders. However, the general closed-loop framework is not just limited to invasive treatment alone. Instead, other, more subtle means of actuation could also be used to close the loop. For instance, music, beverages such as herbal tea and coffee, changes in room lighting could all be used as a means of actuation for regulating arousal. Moreover, cortisol injections at specific times could also be used to close the loop instead of using subcutaneous infusion through an automated pump. The basic closed-loop framework could, therefore, be applicable in a number of other scenarios.

1.3 Challenges and Objectives

Several challenges exist when developing a system-theoretic or state-space formulation governing processes in the human body. A few of them are listed below:

- The time evolution of the underlying state variables and how they are influenced by external factors are often unknown. For instance, a person's mood and emotions vary throughout the day but a precise mathematical equation for how mood or emotions vary is unknown. Moreover, external circumstances impact mood and emotion in different ways. Additionally, the information regarding a particular internal state

within the body may often be distributed among different signals. It is likely, therefore, that there is no single "super feature" which may tell us everything about a particular state.

- Variations exist between individuals. Therefore, any model developed needs to be person-specific. Despite the need for person-specific models, expert labels needed to make use of supervised learning methods are often hard to come by, or are too expensive to obtain. This makes the use of supervised learning more challenging.

- There is a reluctance in the clinical community towards the acceptance of purely black-box methods, and a preference for interpretability. Consider, for instance, the draft guidance document regarding clinical decision support software published by the Food and Drug Administration (FDA) in September 2019 [35]. The document concerns the development of expert software and whether or not it comes under the definition (and regulations) of a device. If expert software is to be excluded from the definition of a medical device and thereby avoid regulations, it should enable a

  > "health care professional to independently review the basis for such recommendations that such software presents so that it is not the intent that such health care professional rely primarily on any of such recommendations to make a clinical diagnosis or treatment decision" [35].

  The document goes on to mention that,

  > "in order to describe the basis for a recommendation, regardless of the complexity of the software and whether or not it is proprietary, the software developer should describe the underlying data used to develop the algorithm and should include plain language descriptions of the logic or rationale used by an algorithm to render a recommendation" [35].

  In other words, the software needs to explain itself. While the guidance document is

only for the purpose of inviting comments, it nonetheless reflects FDA's desire to see interpretability in medical-grade expert systems.

Here we utilize Bayesian methods for estimating internal states within the human body based on sensor observations. We first develop state-space models and corresponding Bayesian filters for estimating sympathetic arousal from skin conductance and heart rate. We secondly develop a similar set of tools to estimate energy from blood cortisol measurements. Machine learning methods are finally used for estimation. Descriptions of these methods are provided in subsequent chapters.

1.4 Approach

Bayesian filtering offers a family of methods for estimating an unknown state variable from observed data. In the case of both skin conductance and cortisol, the signals have a "spikey" appearance to them owing to the mechanism responsible for their generation. Bursts of neural activity to the sweat glands cause variations in the conductivity of the skin. These bursts of neural activity are often modeled as delta functions [36]. Each burst of neuroelectric activity gives rise to a single skin conductance response (SCR) in what is known as the *phasic* component of a skin conductance signal [37]. The SCRs make up the phasic component. This relatively fast-varying phasic component is superimposed on top of a slower-varying *tonic* component [38]. The sequence of SCRs (or equivalently the sequence of underlying neural impulses) forms train of binary observations. We have a similar scenario for cortisol. Cortisol is secreted pulses throughout the day and therefore blood cortisol concentrations also have a "spikey" profile. Fig. 3 depicts a skin conductance signal and a blood cortisol profile. In the figure, the upper sub-panel depicts a skin conductance signal (blue), its tonic component (black) and the neural impulses underlying phasic variations (red). The lower sub-panel depicts blood cortisol measurements taken at 10 min intervals (blue dots), the reconstructed minute-by-minute blood cortisol concentrations (black) and the pulsatile secretions (red).

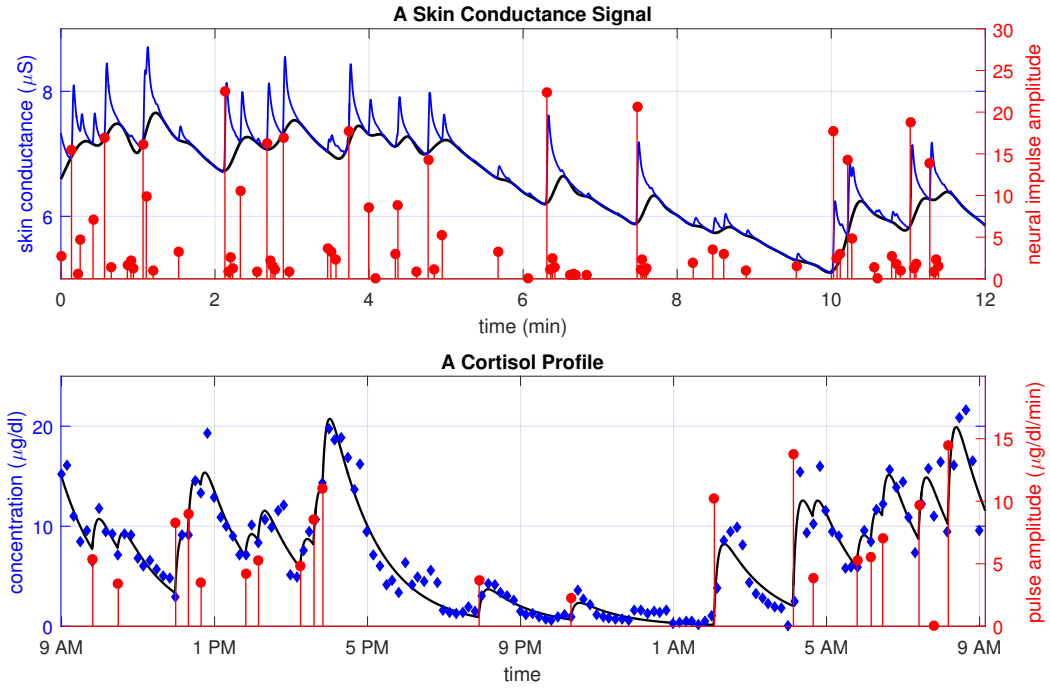

Figure 3: A skin conductance signal and a blood cortisol profile.

A particular class of point process Bayesian filters was pioneered by Professor Emery N. Brown and his Neuroscience Statistics Research Lab of the Massachusetts Institute of Technology and the Massachusetts General Hospital, for estimating unobserved states underlying neural spiking information [39, 40]. These methods have since been adapted and applied to experiments in behavioral learning [41, 42, 43, 44, 45, 46], medically-induced coma [47, 48, 49, 50, 51], anesthetic-induced unconsciousness [52, 53, 54], motor task decoding [55, 56, 57, 58] and sleep studies [59]. The methods relate an unobserved neural state to binary and continuous-valued observations. The filtering itself is typically applied within an expectation-maximization (EM) framework for state estimation and model parameter recovery. We follow the same basic approach when developing filters for estimating sympathetic arousal and energy from skin conductance and heart rate, and blood cortisol measurements respectively. Heart rate can be derived from either an EKG or PPG signal. Fig. 4 depicts part of an EKG signal with its constituent waveforms. R-peaks in

an EKG accompany ventricular contraction (these contractions form a point process) and RR-intervals can be calculated by detecting individual beats.

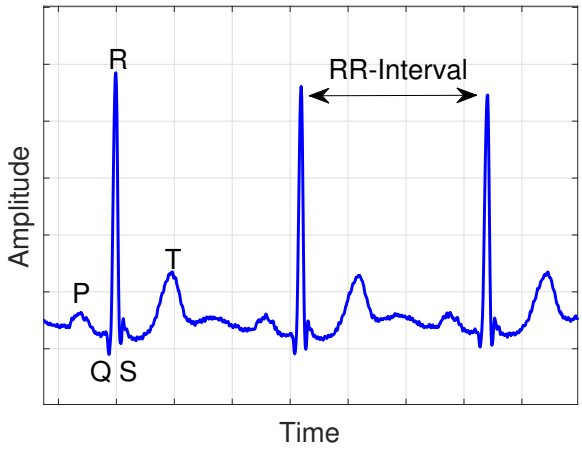

Figure 4: An EKG signal.

1.5 The Need to Estimate an Internal State

Since the methods we develop will likely be embedded within closed-loop control systems, a question may arise as to the need for estimating an internal state when merely the external observation can be controlled. This would be best answered by means of an analogy. Suppose that a person suffers from high blood pressure due to mental stress. Blood pressure alone could be controlled using medication (e.g., medication that removes excess fluid and salt from the body). This, however, would not necessarily solve the deeper issue. An alternate approach would be to try and employ relaxation techniques and/or anxiety medication directed at the brain or nervous system for lowering blood pressure. Our state-space methods are analogous to this second approach where we try to estimate and regulate the *internal* state instead of trying to merely control the *observed symptom*. Note, however, that the analogy to blood pressure strictly does not hold with skin conductance. Lowering blood pressure alone is of clinical importance (e.g., for lowering the risk of stroke). However, merely cooling down the skin if skin conductance is high is not of great clinical significance.

Controlling the external observation instead of the internal energy state is, however, a viable alternative in the case of hormones. For instance, a control mechanism could solely focus on regulating blood cortisol concentrations instead of regulating the internal state. In these situations, however, we do still believe that mathematical models such as the ones we develop help enhance our understanding of how the human body functions. The models may also be able to provide insight into pathological mechanisms in certain disease conditions.

Owing to these reasons, the control mechanism and the means of actuation we envisage are also not primarily directed at the external observation, but rather at the internal state. Sympathetic arousal, for instance, can be controlled by means of actuation such as music or different beverages (e.g., herbal tea, coffee).

1.6 Contributions and Outline

We present several state-space methods in this dissertation for estimating latent variables from physiological signals. In particular, the following contributions are made. We develop:

- Point process state-space methods to estimate sympathetic arousal from skin conductance measurements
- Multi-timescale point process state-space methods to estimate the same from a combination of skin conductance and heart rate measurements
- Point process state-space methods to estimate energy production from pulsatile blood cortisol measurements
- A hybrid machine learning-based state estimation method for marked point process data (MPP) that is able to combine physiological data and an external influence

The rest of this dissertation is organized as follows. Chapter 2 describes the Bayesian filters and EM algorithms used for estimating sympathetic arousal from skin conductance. These methods are extended in chapter 3 to include heart rate. Chapter 4 describes approaches to estimate energy from blood cortisol measurements. In chapter 5, we describe

how machine learning methods can be used for state estimation in both skin conductance and cortisol-based applications. We finally conclude with chapter 6 summarizing our work and noting several possible directions of future research.

2 Sympathetic Arousal Estimation using Skin Conductance Data

2.1 Sweat Secretion, Skin Conductance and Arousal

In this chapter, we describe different approaches for estimating arousal from skin conductance measurements. Nerve fibers from the sympathetic branch of the autonomic nervous system innervate the sweat glands [60]. Consequently, a skin conductance signal, which varies based on sweat secretions, functions as an index of sympathetic arousal [61]. Sympathetic innervation of the sweat glands is primarily cholinergic and involves the neurotransmitter acetylcholine [60]. Several different skin conductance features have been used in the literature as markers of arousal. Firstly, the rate at which SCRs occur has been taken as an index of arousal—the higher the arousal, the higher the rate of SCR occurrence. SCR rate has been used as a marker of arousal in experiments involving thought suppression [62], alcohol craving [63] and audio processing [64]. Secondly, SCR amplitude has been considered an index of arousal. This has been used in studies involving emotional visual stimuli [65] and sounds [66]. Finally, the tonic level has been used as an index of arousal in experiments involving biofeedback tasks [67], antisocial behavior [68] and the presentation of visual stimuli [69]. SCR rate, SCR amplitude and the tonic level are the three most commonly reported skin conductance markers of autonomic arousal in the literature [70].

We describe three different methods for skin conductance-based arousal estimation in this chapter. The methods use different combinations of the three features just referred to. The methods are:

- Estimation based on one binary observation
- Estimation based on one binary and two continuous observations
- Estimation based on MPP observations

Estimation based on one binary observation followed the methodology in [41] and does

not represent an original method we developed. Instead it was our our earliest attempt at estimating arousal from skin conductance and only utilized the SCR occurrences (or equivalently the neural impulses that underlie the SCRs) [71, 24]. We next developed an estimation method based on one binary and two continuous observations. This represents the most complicated case and is described in [72]. The method utilizes the SCR occurrences (binary feature) and a phasic-derived value and the tonic component (two continuous features) for estimation. The MPP approach was our third attempt at estimating arousal [73]. It too considered the neural impulses underlying phasic variations, but took into account the impulse amplitudes as well. In this chapter, we abbreviate the first approach as the binary observation-based state estimation (BOBSE) method, the second as the binary and continuous observation-based state estimation (BCOBSE) method and the third as the MPP-based state estimation (MPPSE) method. From a purely mathematical standpoint, the derivation of the Bayesian filter equations for the BCOBSE method is the most complicated. By dropping a single variable in the derivation, we can obtain part of the equations for MPPSE. Dropping yet another variable in the derivation yields the equations for BOBSE. We will therefore describe the state-space model used in BCOBSE first in the Methods section, and thereafter describe the mathematically simpler models used in BOBSE and MPPSE.

2.2 Data

2.2.1 Neurological Status Assessment Dataset

We first use the Non-EEG Biosignals Dataset for Assessment and Visualization of Neurological Status (which we abbreviate as the Neurological Status Assessment Dataset) [74]. This contains skin conductance recordings from a group of college students who were exposed to several conditions meant to elicit stress. The experiment included physical, cognitive and emotional stressors. Subjects were required to stand, walk and then walk/jog during the physical stress period. This lasted for approximately 5 min. Here we discard

the physical stress portion and only consider remaining the psychological components of the experiment. The cognitive stress period consisted of two different tasks. In the first task, the subjects had to count backwards in 7's beginning at 2485. This lasted for approximately 3 min. The Stroop test was the second cognitive stressor. Different variants of the Stroop color-word association test appear in the literature. In this particular experiment, the names of colors written in different-colored ink were shown to the subjects. The subjects had to first read out the color name and then state the color of the ink [74]. The need to perform this correctly requires concentration on the part of the subjects and is known to generate psychological stress [75]. The Stroop test lasted for approximately 2 min. A buzzer notified the subjects of errors. Subjects were shown a horror movie clip for 5 min as an emotional stressor. The physical, cognitive and emotional stress periods were interspersed by 5 min intervals of relaxation. The authors of the dataset also categorized the 40 s just prior to the cognitive stressors as a stress period as they noted the subjects visibly displaying signs of stress while even being given the instructions [74].

2.2.2 Driver Stress Dataset

We secondly use the Stress Recognition in Automobile Drivers Dataset (we abbreviate this as the Driver Stress Dataset) [76]. This contains skin conductance recordings from subjects who drove along a set route in Boston as part of the experiment. The route consisted of toll roads, highways and city driving. As we pointed out earlier in [71], the precise timings of the different road conditions are unavailable and we used the data from a single subject for whom the road condition timings had to be approximately matched to a figure in [76] which originally described the dataset.

2.3 Motion Artifact Contamination

Data preprocessing typically involves downsampling to 4 Hz, lowpass filtering at 0.5 Hz and tonic-phasic separation using cvxEDA [77]. However, prior to lowpass filtering, we

usually employ a motion artifact cleaning step. Here we visually inspect the data, identify contaminated segments and interpolate over them. Motion artifacts can be a major source of contamination, especially with data acquired from wearable devices in unconstrained environments. These artifacts can also occur in signals acquired from research-grade equipment in laboratory settings, though not to an extent as great. While we currently clean motion artifacts manually, adaptive signal processing methods can also be used to suppress this type of nonlinear contamination. Another alternative to purely using signal processing methods for motion artifact suppression is to develop flexible or adhesive sensors (e.g., drawn-on-skin electronics) so that the movement of the sensors relative to the skin is kept minimal [78]. Adaptive filters could then be added on as another layer for further noise suppression.

2.4 Methods

We first describe the full model incorporating one binary and two continuous observations (i.e., BCOBSE). Mathematically, this state-space model gives rise to the most complicated filter equations. The filter equations for the other models can be easily derived by dropping the continuous variables one at a time.

2.4.1 State-space Model with One Binary and Two Continuous Observations

Random walks and first-order autoregressive models have frequently been used to capture the time evolution of neural states that cannot be directly observed (e.g., learning [41, 44, 46, 79], sleep [59] and neural states underlying spiking activity [39]). Similar to [39], we assume that sympathetic arousal x_k evolves with time as

$$x_k = \rho x_{k-1} + \alpha I_k + \varepsilon_k, \tag{1}$$

where $\varepsilon_k \sim \mathcal{N}(0, \sigma_\varepsilon^2)$ is process noise and I_k is an indicator function representing external stimuli. ρ and α are coefficients to be determined. We consider three different skin conductance features to estimate x_k here.

We first consider the occurrence of SCRs. SCRs can be detected as phasic peaks that exceed a threshold between 0.01–0.05 µS [38]. We assign $m_k = 1$ or $m_k = 0$ based on whether or not an SCR occurred at the k^{th} time index using a threshold of 0.015 µS. The occurrence of SCRs follows a Bernoulli distribution with a density function $p_k^{m_k}(1-p_k)^{1-m_k}$ where p_k is the probability that $m_k = 1$. We relate sympathetic arousal to the occurrence of SCRs using the theory of generalized linear models. We use a logit transformation following the suggestion in [80]. This yields

$$\log\left(\frac{p_k}{1-p_k}\right) = \beta_0 + \beta_1 x_k \implies p_k = \frac{1}{1+e^{-(\beta_0+\beta_1 x_k)}}, \tag{2}$$

where β_0 and β_1 are constant coefficients. Therefore,

$$P(m_k|x_k) = p_k^{m_k}(1-p_k)^{1-m_k} = \left[\frac{1}{1+e^{-(\beta_0+\beta_1 x_k)}}\right]^{m_k}\left[\frac{e^{-(\beta_0+\beta_1 x_k)}}{1+e^{-(\beta_0+\beta_1 x_k)}}\right]^{1-m_k}. \tag{3}$$

Here we set $\beta_1 = 1$ and calculate β_0 empirically. Assuming that $x_k \approx 0$ at the very beginning, we have

$$\log\left(\frac{p_0}{1-p_0}\right) \approx \beta_0. \tag{4}$$

We can approximate p_0 by the baseline probability of an SCR occurring during the entire experiment.

Secondly, we consider the continuous-valued tonic level s_k which is also known to be related to arousal [81]. Other neural state estimation methods (e.g., [44, 46]) have assumed linear relationships between continuous-valued observations and the latent state to

be determined. We too take s_k to be linearly related to x_k as

$$s_k = \delta_0 + \delta_1 x_k + w_k, \tag{5}$$

where $w_k \sim \mathcal{N}(0, \sigma_w^2)$ captures sensor noise and modeling error, and δ_0 and δ_1 are regression coefficients to be determined.

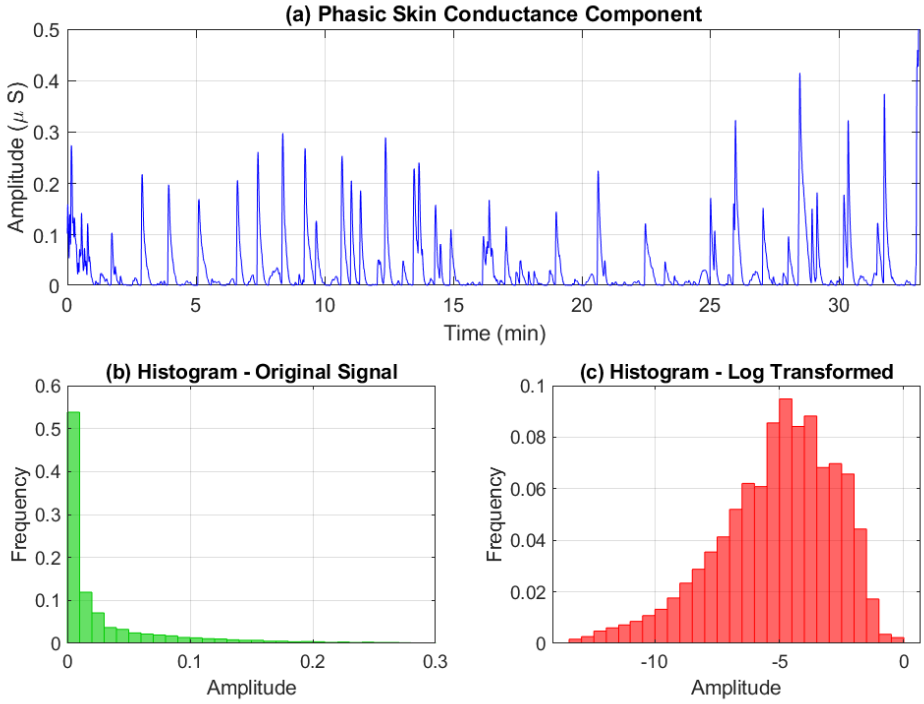

Figure 5: Phasic skin conductance skew correction via log transformation.

We next consider the rapidly-fluctuating phasic component $\tilde{r}_k$. Two main aspects are to be noted regarding the phasic component. Firstly, its distribution is skewed. A logarithmic or square-root transformation is commonly suggested to correct skew in skin conductance features [82]. Here we apply a logarithmic transformation (Fig. 5). Secondly, it is the amplitudes of the SCRs that are considered to be related to sympathetic arousal [83]. Therefore, we derive an artificial signal r_k by interpolating over the SCR peaks and the first and last values of the log-transformed $\tilde{r}_k$. Taking $r^* = \{\tilde{r}_1, \tilde{r}_K\} \cup \{\tilde{r}_k | m_k = 1\}$ to denote

the phasic SCR peaks along with the first and last values, r_k is derived by applying a cubic interpolation over $\log r^*$ (Fig. 6). The positive skew of the SCR amplitudes is one that has been noted in the literature and the logarithmic transformation is often used for correction [84]. Similar to the case of s_k, we assume a linear relationship between r_k and x_k as well. Therefore,

$$r_k = \gamma_0 + \gamma_1 x_k + v_k, \tag{6}$$

where the coefficients γ_0 and γ_1 are similar to δ_0 and δ_1, and $v_k \sim \mathcal{N}(0, \sigma_v^2)$ represents a noise term similar to w_k. $\mathcal{M}^K = \{m_1, m_2, \ldots, m_K\}$, $\mathcal{R}^K = \{r_1, r_2, \ldots, r_K\}$ and $\mathcal{S}^K = \{s_1, s_2, \ldots, s_K\}$ form the complete series of skin conductance observations (Fig. 7). In the figure, the sub-panels respectively depict: (a) the skin conductance signal z_k; (b) the phasic component $\tilde{r}_k$ (the red circles depict the SCR peaks); (c) the phasic-derived signal r_k; (d) the tonic component s_k.

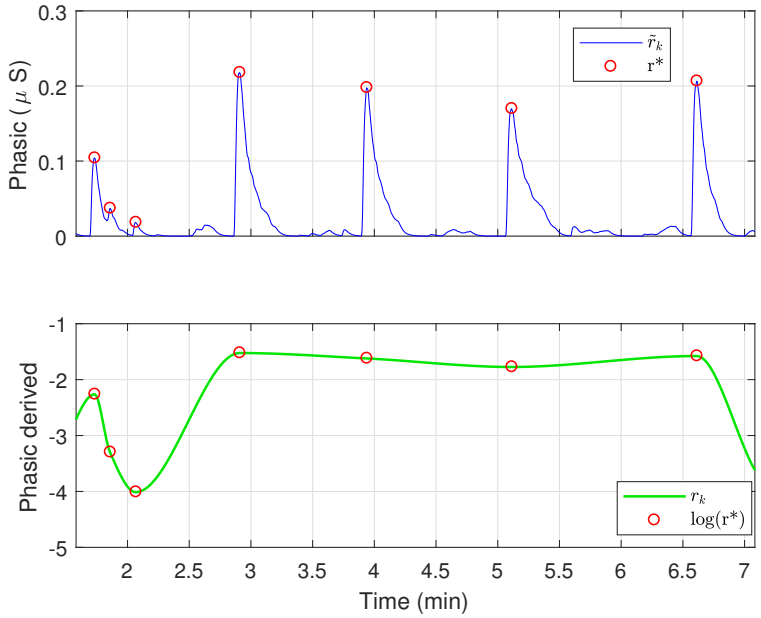

Figure 6: Extraction of the phasic-derived component r_k from $\tilde{r}_k$.

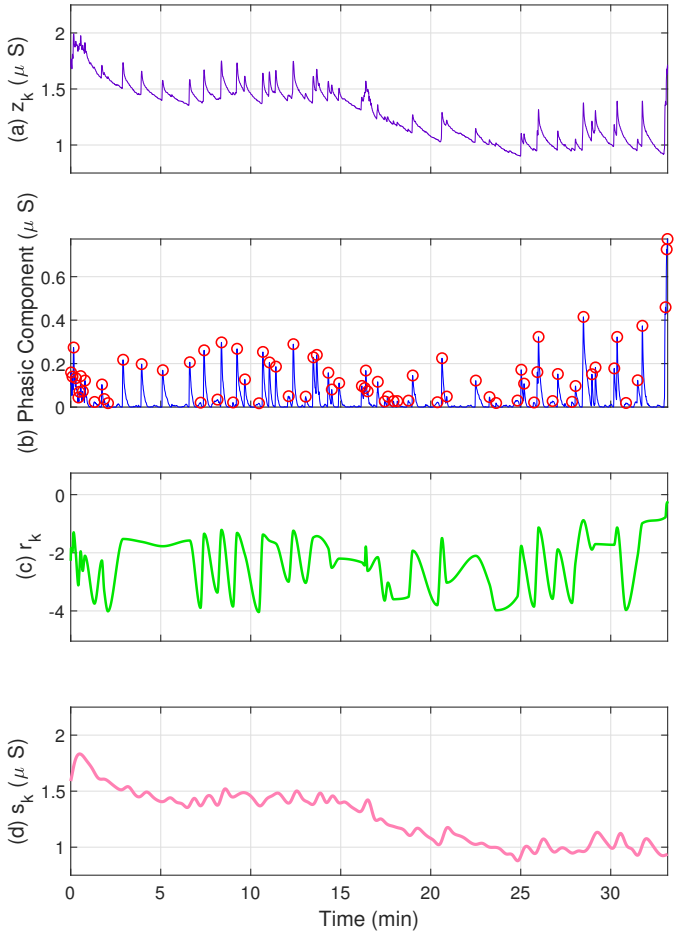

Figure 7: Constituent components of a skin conductance signal.

Given $\mathcal{M}^K$, $\mathcal{R}^K$, $\mathcal{S}^K$, and letting $\mathcal{Y}^K = \{\mathcal{M}^K, \mathcal{R}^K, \mathcal{S}^K\}$, we wish to determine $\mathcal{X}^K = \{x_1, x_2, \ldots, x_K\}$. To do so, we use Bayesian filtering within an EM framework. In the E-step, we first obtain the estimates of $\mathcal{X}^K$ conditioned on $\mathcal{Y}^K$ and the model parameters $\Theta = \{\alpha, \rho, \gamma_0, \gamma_1, \delta_0, \delta_1, \sigma_v^2, \sigma_w^2, \sigma_\varepsilon^2\}$. In the M-step, we choose the model parameters that maximize the expected value of the complete data log-likelihood. The algorithm alternates between the E-step and the M-step until convergence.

We estimate $\mathcal{X}^K$ in two steps at the E-step using both a forward filter and a backward smoother. We first make a Gaussian approximation to the posterior distribution $p(x_k|y^k)$

similar to [43, 46] in order to obtain the following filter equations:

$$x_{k|k-1} = \rho x_{k-1|k-1} + \alpha I_k, \tag{7}$$

$$\sigma^2_{k|k-1} = \rho^2 \sigma^2_{k-1|k-1} + \sigma^2_{\varepsilon}, \tag{8}$$

$$C_k = \frac{\sigma^2_{k|k-1}}{\sigma^2_v \sigma^2_w + \sigma^2_{k|k-1}(\gamma_1^2 \sigma^2_w + \delta_1^2 \sigma^2_v)}, \tag{9}$$

$$\begin{aligned} x_{k|k} = x_{k|k-1} + C_k \Big[& \sigma^2_v \sigma^2_w (m_k - p_{k|k}) + \gamma_1 \sigma^2_w (r_k - \gamma_0 - \gamma_1 x_{k|k-1}) \\ & + \delta_1 \sigma^2_v (s_k - \delta_0 - \delta_1 x_{k|k-1}) \Big], \end{aligned} \tag{10}$$

and

$$\sigma^2_{k|k} = \left[\frac{1}{\sigma^2_{k|k-1}} + p_{k|k}(1 - p_{k|k}) + \frac{\gamma_1^2}{\sigma^2_v} + \frac{\delta_1^2}{\sigma^2_w} \right]^{-1}. \tag{11}$$

Now,

$$p_{k|k} = \frac{1}{1 + e^{-(\beta_0 + x_{k|k})}} \tag{12}$$

causes $x_{k|k}$ to appear on both sides of eq. (10) and has to be solved numerically via the Newton-Raphson method [41]. Similar to a Kalman filter update step, eq. (10) contains innovation terms for the continuous variables r_k and s_k and also a term comparing the binary value m_k to its predicted probability $p_{k|k}$. The smoothed estimates $x_{k|K}$ and $\sigma^2_{k|K}$ conditioned on having observed all the data up to time index K [41, 85] are obtained as follows:

$$A_k \triangleq \rho \frac{\sigma^2_{k|k}}{\sigma^2_{k+1|k}}, \tag{13}$$

$$x_{k|K} = x_{k|k} + A_k (x_{k+1|K} - x_{k+1|k}), \tag{14}$$

and

$$\sigma_{k|K}^2 = \sigma_{k|k}^2 + A_k^2(\sigma_{k+1|K}^2 - \sigma_{k+1|k}^2). \tag{15}$$

The Θ parameters are determined in the M-step. These parameters are selected to maximize the expected value of the complete data log-likelihood. We make use of the state-space covariance algorithm when deriving the parameter estimates [86]. Defining

$$U_k = x_{k|K}^2 + \sigma_{k|K}^2 \tag{16}$$

and

$$U_{k,k+1} = x_{k|K}x_{k+1|K} + A_k\sigma_{k+1|K}^2, \tag{17}$$

we obtain the following parameter estimates for the $(l+1)^{\text{th}}$ EM iteration:

$$\begin{bmatrix} \rho^{(l+1)} \\ \alpha^{(l+1)} \end{bmatrix} = \begin{bmatrix} \sum_{k=1}^{K-1} U_k & \sum_{k=2}^{K} I_k x_{k-1|K} \\ \sum_{k=2}^{K} I_k x_{k-1|K} & \sum_{k=1}^{K} I_k^2 \end{bmatrix}^{-1} \times \begin{bmatrix} \sum_{k=1}^{K-1} U_{k,k+1} \\ \sum_{k=2}^{K} I_k x_{k|K} \end{bmatrix}, \tag{18}$$

$$\begin{bmatrix} \gamma_0^{(l+1)} \\ \gamma_1^{(l+1)} \end{bmatrix} = \begin{bmatrix} K & \sum_{k=1}^{K} x_{k|K} \\ \sum_{k=1}^{K} x_{k|K} & \sum_{k=1}^{K} U_k \end{bmatrix}^{-1} \times \begin{bmatrix} \sum_{k=1}^{K} r_k \\ \sum_{k=1}^{K} r_k x_{k|K} \end{bmatrix}, \tag{19}$$

$$\begin{aligned} \sigma_v^{2(l+1)} = \frac{1}{K}\Bigg[& \sum_{k=1}^{K} r_k^2 + K\gamma_0^{2(l+1)} + \gamma_1^{2(l+1)} \sum_{k=1}^{K} U_k - 2\gamma_0^{(l+1)} \sum_{k=1}^{K} r_k - 2\gamma_1^{(l+1)} \sum_{k=1}^{K} x_{k|K} r_k \\ & + 2\gamma_0^{(l+1)}\gamma_1^{(l+1)} \sum_{k=1}^{K} x_{k|K} \Bigg], \end{aligned} \tag{20}$$

and

$$\begin{aligned} \sigma_\varepsilon^{2(l+1)} = \frac{1}{K}\Bigg[& \sum_{k=2}^{K} U_k - 2\rho^{(l+1)} \sum_{k=1}^{K-1} U_{k,k+1} + \rho^{2(l+1)} \sum_{k=1}^{K-1} U_k - 2\alpha^{(l+1)} \sum_{k=2}^{K} I_k x_{k|K} \\ & + 2\alpha^{(n+1)}\rho^{(l+1)} \sum_{k=2}^{K} I_k x_{k-1|K} + \alpha^{2(l+1)} \sum_{k=1}^{K} I_k^2 \Bigg]. \end{aligned} \tag{21}$$

The M-step updates for $\delta_0^{(l+1)}$ and $\delta_1^{(l+1)}$ may be obtained by replacing r_k with s_k in eq. (19). Likewise, $\sigma_w^{2(l+1)}$ can be obtained similar to eq. (20).

Note that in this particular state-space model, the posterior $p(x_k|y^k)$ contains the product of the three terms $P(m_k|x_k)$, $p(r_k|x_k)$ and $p(s_k|x_k)$ corresponding to the binary observation and the two continuous observations.

2.4.2 State-space Model with One Binary Observation

If the two continuous variables r_k and s_k are removed from the state-space model, $p(x_k|y^k)$ will no longer contain $p(r_k|x_k)$ and $p(s_k|x_k)$, but will only contain $P(m_k|x_k)$. In this case, we obtain the filter equations in [41] which were originally developed to estimate the cognitive learning state of an animal engaged in a experiment comprising of a sequence of trials with correct/incorrect responses. Here we adapt the model and consider the SCR occurrences or the neural impulses underlying the SCRs to form the binary observations m_k. Assuming a random walk

$$x_k = x_{k-1} + \varepsilon_k, \tag{22}$$

for the state equation, the filter equations are as follows:

$$x_{k|k-1} = x_{k-1|k-1}, \tag{23}$$

$$\sigma^2_{k|k-1} = \sigma^2_{k-1|k-1} + \sigma^2_{\varepsilon}, \tag{24}$$

$$x_{k|k} = x_{k|k-1} + \sigma^2_{k|k-1}(m_k - p_{k|k}), \tag{25}$$

and

$$\sigma^2_{k|k} = \left[\frac{1}{\sigma^2_{k|k-1}} + p_{k|k}(1 - p_{k|k})\right]^{-1}. \tag{26}$$

In this case as well, $p_{k|k}$ appears on both sides of the state update equation in (25) and has to be solved using Newton's method. The smoother equations do not change except for the absence of ρ. We also make a small change with this particular state-space model in that we estimate the initial state x_0. Therefore we have the following M-step updates:

$$\begin{aligned}\sigma_{\varepsilon}^{2(l+1)} =& \frac{2}{K+1}\left[\sum_{k=2}^{K}(\sigma_{k|K}^2 + x_{k|K}^2) - \sum_{k=2}^{K}(A_k\sigma_{k|K}^2 + x_{k|K}x_{k-1|K})\right] \\ &+ \frac{1}{K+1}\left[\frac{3}{2}x_{1|K}^2 + 2\sigma_{1|K}^2 - (\sigma_{K|K}^2 + x_{K|K}^2)\right] \end{aligned} \tag{27}$$

and

$$x_0^{(l+1)} = \frac{1}{2}x_{1|K}. \tag{28}$$

2.4.3 State-space Model with One Marked Point Process Observation

Our third and final model is an extension of the previous approach and accounts for the neural impulse amplitudes in addition to their occurrence alone. In this model as well, we assume that x_k evolves with time following a random walk and that x_k is related to the probability of neural impulse occurrence p_k via a sigmoid. Moreover, similar to eq. (6), we also take the amplitudes of the impulses r_k to be linearly related to x_k as

$$r_k = \gamma_0 + \gamma_1 x_k + v_k, \tag{29}$$

where γ_0 and γ_1 have their usual meanings and $v_k \sim \mathcal{N}(0, \sigma_v^2)$. Consequently, the joint density function for the observed neural stimuli is

$$p(m_k \cap r_k | x_k) = \begin{cases} 1 - p_k & \text{if } m_k = 0 \\ p_k \frac{1}{\sqrt{2\pi\sigma_v^2}} e^{\frac{-(r_k - \gamma_0 - \gamma_1 x_k)^2}{2\sigma_v^2}} & \text{if } m_k = 1. \end{cases} \tag{30}$$

Accordingly, an amplitude is absent if there is no impulse and the amplitude is Gaussian-distributed when an impulse is present. The filter equations for estimating x_k are as follows:

$$x_{k|k-1} = x_{k-1|k-1} \tag{31}$$

and

$$\sigma^2_{k|k-1} = \sigma^2_{k-1|k-1} + \sigma^2_{\varepsilon}. \tag{32}$$

If $m_k = 0$, we have

$$x_{k|k} = x_{k|k-1} + \sigma^2_{k|k-1}(m_k - p_{k|k}), \tag{33}$$

$$\sigma^2_{k|k} = \left[\frac{1}{\sigma^2_{k|k-1}} + p_{k|k}(1 - p_{k|k})\right]^{-1}, \tag{34}$$

and if $m_k = 1$, we instead have

$$x_{k|k} = x_{k|k-1} + \frac{\sigma^2_{k|k-1}}{\gamma_1^2\sigma^2_{k|k-1} + \sigma^2_v}[\sigma^2_v(m_k - p_{k|k}) + \gamma_1(r_k - \gamma_0 - \gamma_1 x_{k|k-1})], \tag{35}$$

and

$$\sigma^2_{k|k} = \left[\frac{1}{\sigma^2_{k|k-1}} + p_{k|k}(1 - p_{k|k}) + \frac{\gamma_1^2}{\sigma^2_v}\right]^{-1}. \tag{36}$$

It is of interest to note that the filter equations switch between those in [41] (estimation based on one binary variable) and [44] (estimation based on one binary variable and one continuous variable) depending on whether $m_k = 0$ or $m_k = 1$. In [41], Smith et al., developed a state-space method to estimate an unobserved cognitive learning state from binary correct/incorrect responses in behavioral experiments involving animal models. Prerau et al., [44] extended the model to incorporate reaction times (i.e., a continuous variable) so

that the estimated learning state was based on both a binary variable and a continuous variable.

The M-step update equations are as follows:

$$\begin{bmatrix} \gamma_0^{(l+1)} \\ \gamma_1^{(l+1)} \end{bmatrix} = \begin{bmatrix} |\tilde{K}| & \sum_{k\in\tilde{K}} x_{k|K}, \\ \sum_{k\in\tilde{K}} x_{k|K} & \sum_{k\in\tilde{K}} U_k \end{bmatrix}^{-1} \times \begin{bmatrix} \sum_{k\in\tilde{K}} r_k \\ \sum_{k\in\tilde{K}} r_k x_{k|K} \end{bmatrix}, \tag{37}$$

$$\begin{aligned} \sigma_v^{2(l+1)} = \frac{1}{|\tilde{K}|}\Bigg[& \sum_{k\in\tilde{K}} r_k^2 + |\tilde{K}|\gamma_0^{2(l+1)} + \gamma_1^{2(l+1)} \sum_{k\in\tilde{K}} U_k - 2\gamma_0^{(l+1)} \sum_{k\in\tilde{K}} r_k \\ & - 2\gamma_1^{(l+1)} \sum_{k\in\tilde{K}} x_{k|K} r_k + 2\gamma_0^{(l+1)}\gamma_1^{(l+1)} \sum_{k\in\tilde{K}} x_{k|K} \Bigg], \end{aligned} \tag{38}$$

and

$$\sigma_\varepsilon^{2(l+1)} = \frac{1}{K}\left[\sum_{k=2}^{K} U_k - 2\sum_{k=1}^{K-1} U_{k,k+1} + \sum_{k=1}^{K-1} U_k \right], \tag{39}$$

where $\tilde{K}$ represents the indices at which $m_k = 1$.

2.5 Results and Discussion

2.5.1 Simulated Data

Two of the three EM algorithms described above are original formulations (the BCOBSE and MPPSE methods). We first evaluated the performance of both of them on simulated data.

We first consider the BCOBSE method. To verify the performance of the corresponding EM algorithm, we simulated data for an arousal state x_k using the parameters $\Theta = \{0.04, 0.995, -4.595, 0.35, 0.4, -0.7, 0.2, 0.002, 0.005, 0.03\}$. These values were chosen based on prior experience with experimental data. We arbitrarily set $I_k = 1$ at 25 time instances and 0 elsewhere. The EM stopping criteria is similar to [41, 44] and we consider the parameters to have converged once their absolute mean deviation in consecutive iterations does

not exceed 10^{-8}.

Recall that we calculate $\beta_0 = \log[p_0(1-p_0)^{-1}]$ using an approximation for p_0 as the average probability that $m_k = 1$ [71]. We label this approximate probability $\hat{p}_0$. We set β_0 corresponding to a true $p_0 = 0.01$ for simulating data, i.e., $\beta_0 = \log[0.01(1-0.01)^{-1}]$. We next generated two synthetic datasets for which the approximation $\hat{p}_0$ was slightly lower than, and then slightly higher than 0.01. The model parameters recovered by the EM algorithm $\hat{\Theta}$ are shown in Table 1. The state estimates and fits to one set of simulated data are shown in Fig. 8. In the figure, the sub-panels respectively depict: (a) the probability of binary event occurrence p_k (blue) and its estimate (red, the green and black dots on the upper part of the sub-panel denote the presence or absence of the binary events); (b) the first continuous-valued variable r_k (blue) and its estimate (red); (c) the second continuous-valued variable s_k (blue) and its estimate (red); (d) the state x_k (blue) and its estimate (red, the cyan and black dots on the lower part of the sub-panel denote the presence or absence of binary inputs); (e) the quantile-quantile (QQ) plot for the residual errors of x_k. While the proposed scheme can recover model parameters and estimate x_k reasonably well, the less-than-true approximation ($\hat{p}_0 < 0.01$) can cause x_k to slightly overestimate the true state and the higher-than-true approximation ($\hat{p}_0 > 0.01$) can cause x_k to slightly underestimate the true state.

Validation with simulated data was similar for the MPPSE method. Again, we generated two datasets with a true $p_0 = 0.05$, but where the approximations were slightly higher than and slightly lower than 0.05. The model parameters and their estimates for both cases are shown in Table 2. Fig. 9 shows the state estimation results in one of the cases. In the figure, the sub-panels respectively depict: (a) the MPP observations (blue) and the estimated fit to r_k (red); (b) the probability of point process event occurrence p_k (blue) and its estimate (red); (c) the state x_k (blue) and its estimate (red); (d) the QQ plot for the residual errors of x_k. In general, there is a good fit to the simulated data although the state estimate can deviate from the true value in regions where there is no impulse for a long period of time.

Table 1: Parameter estimation with simulated data – one binary and two continuous observations

True value	Estimate ($\hat{p}_0 < 0.01$)	Estimate ($\hat{p}_0 > 0.01$)
$\alpha = 0.04$	0.018	0.021
$\rho = 0.995$	0.994	0.994
$\beta_0 = -4.595$	-4.64	-4.057
$\beta_1 = 1$	1 (set)	1 (set)
$\gamma_0 = 0.35$	0.233	0.608
$\gamma_1 = 0.4$	0.450	0.377
$\delta_0 = -0.7$	-0.758	-0.571
$\delta_1 = 0.2$	0.224	0.187
$\sigma_V^2 = 0.002$	0.0022	0.0016
$\sigma_W^2 = 0.005$	0.0046	0.0050
$\sigma_\varepsilon^2 = 0.03$	0.0219	0.0216

The estimates are closer to the true state in regions where more impulses tend to occur.

Table 2: Parameter estimation with simulated data – one MPP observation

True value	Estimate ($\hat{p}_0 > 0.05$)	Estimate ($\hat{p}_0 < 0.05$)
$\gamma_0 = 0.2$	0.27341	0.17119
$\gamma_1 = 0.7$	0.68694	0.67415
$\sigma_v^2 = 0.05$	0.04840	0.04885
$\sigma_\varepsilon^2 = 0.005$	0.00417	0.00366

For both the BSOBSE and MPPSE methods, the QQ plots are close to the 45° diagonal indicating a Gaussian distribution of the residual errors. Coupled together with the model parameter estimates and the fits to the true observations and states, the results indicate the ability of our methods to recover x_k accurately from a given set of observations.

A further issue is to be pointed out with regard to the MPPSE method. In regions where an impulse is absent for a long time (due to x_k being low, and consequently the impulse occurrence probability being small), the estimate can fail to capture small deviations in the true sympathetic arousal state. Since estimation is performed by only observing $m_k = 0$ when neural impulses are absent, the algorithm has limited information to work with at these locations. Consequently, if an impulse is absent for a long period of time, then the EM estimate is unable to capture small variations in x_k during that time. In the real-world, this

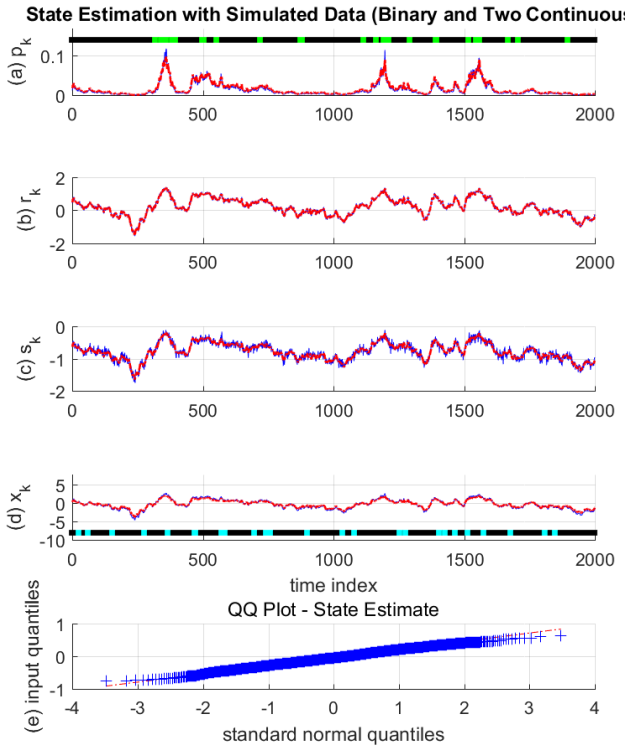

Figure 8: State estimation on simulated data for model with one binary and two continuous observations.

would translate to scenarios where there are small variations in an individual's emotional arousal, but none so large as to cause noticeable SCRs. If no SCRs occur in this manner (or no underlying neural impulses occur), then the MPPSE method is unable to capture the tiny emotional variations during that time. This issue would, however, be common to other skin conductance-based approaches for emotion recognition as well—if no SCRs occur, then the algorithms have little to perform estimation on.

Note further that the EM algorithms place no constraints on the signs of r_k or s_k (i.e., to be either positive or negative). Therefore, the algorithms can be utilized in more general applications as well. Skin conductance represents a special case where some of the observations can only be positive.

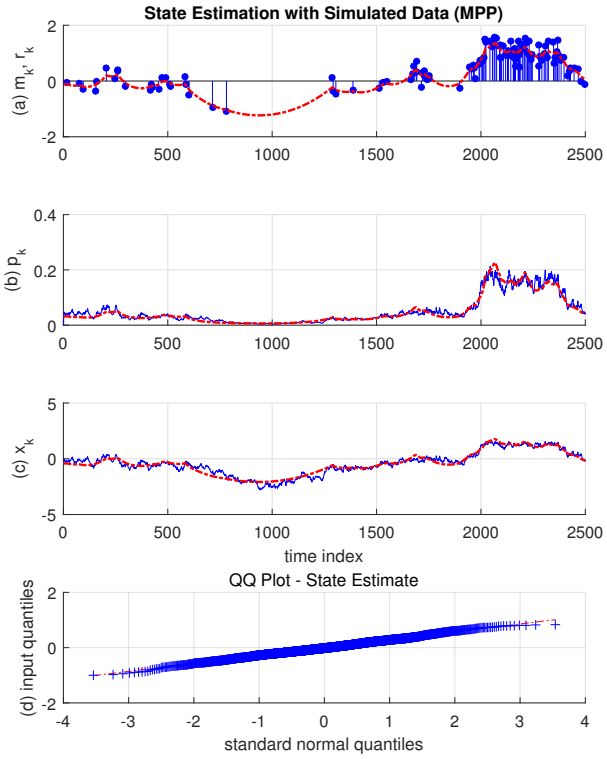

Figure 9: State estimation on simulated data for model with one MPP observation.

2.5.2 Experimental Data

On experimental data, we compare the performance of the BCOBSE, BOBSE and MPPSE methods on representative cases of the two datasets described earlier. We used deconvolved skin conductance data for BOBSE and MPPSE [24, 73] and SCR peak detection for BCOBSE [72]. In each case, we pay particular attention to the state estimate x_k and the high arousal index (HAI) calculated as $p(x_k > x_{\text{threshold}})$. The HAI is inspired by the term known as the ideal observer certainty level in [41]. It originally expressed the probability that a point process event occurred by more than just chance in a behavioral learning experiment. Here it captures the probability that binary events occur above their natural baseline rate. Since p_k is related to x_k according to (2), the certainty level can also be calculated based on the probability that the state x_k exceeds an equivalent threshold.

The threshold would capture the physiological baseline arousal state value and would in turn be related to the baseline impulse occurrence probability. We set $x_{\text{threshold}}$ to the median state value considering that the experiments in the Neurological Status Assessment Dataset and the Driver Stress Dataset contain both relaxation and stress periods [24].

Unlike in the case of simulated data, we place additional constraints when performing state estimation on experimental data with the BCOBSE method. Here there is a tendency for the model parameters to converge to a location where there is an almost perfect fit to one of the continuous-valued observations (either r_k or s_k). It is likely that local extrema exist in the model parameter search space at these points and the EM algorithm can converge to them. In order to prevent x_k from overfitting to r_k or s_k, we first divide r_k and s_k by their respective standard deviations and carefully monitor the variance terms $\sigma_v^{2(l+1)}$ and $\sigma_w^{2(l+1)}$ at each iteration. The model parameters related to r_k and s_k are only allowed to update if the absolute difference between the two variance terms exceeds a threshold (this prevents one of the variance terms being driven down at the expense of the other, which is what happens during overfitting). This largely helps control the degree of overfitting to either r_k or s_k. Fig. 10 shows the arousal estimation results using the BCOBSE method (the external input I_k was not used here). In both sub-figures, the sub-panels respectively depict: (a) the skin conductance signal (the green and black dots above it depict the detected SCR peak locations); (b) the phasic-derived value (solid green line) and its fit (dotted line); (c) the tonic level (solid light mauve line) and its fit (dotted line); (d) the arousal state estimate and its 95% confidence limits; (e) the probability of SCR occurrence and its 95% confidence limits; (f) HAI. The regions above 90% and below 10% are shaded in red and green respectively in sub-panel (f). In the left sub-figure, the background colors in turn denote the instruction period (red), the counting task (green), the Stroop test (orange), relaxation (blue) and emotional stress (yellow – horror movie clip) respectively. In the right sub-figure, the background colors denote rest (green), city driving (blue), toll roads (red) and highways (yellow). On the Neurological Status Assessment Dataset, the subject's

HAI is above 90% during the initial cognitive stressors and extends considerably into the relaxation interval. Also, there is a sizeable increase that even exceeds 90% at the beginning of the emotional stress period. On the Driver Stress Dataset, the BCOBSE method provides a somewhat jagged arousal state estimate. The HAI also has sharp unintuitive rises and falls indicating an almost binary-like transition. There is also no consistent variation in arousal depending on the road type. It is also to be noted that the large increase in HAI during the second rest period was due to the subject being agitated because of a need to use the restroom [76].

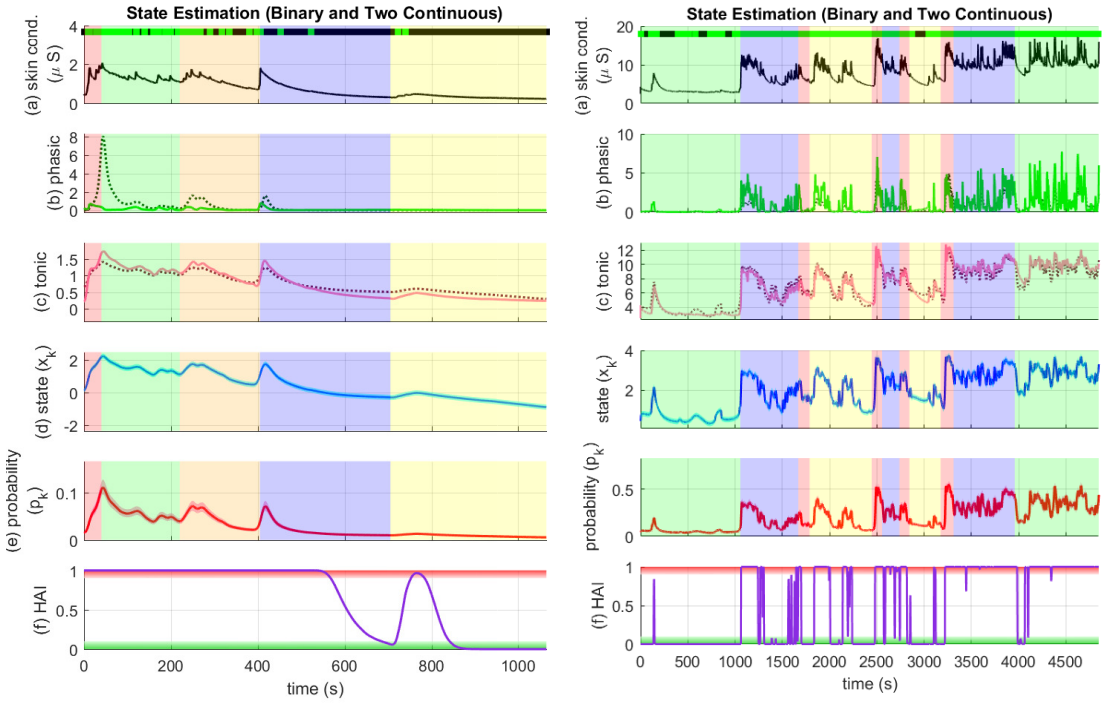

Figure 10: Sympathetic arousal estimation using model with one binary and two continuous observations.

Fig. 11 shows the state estimates using the BOBSE method. In both sub-figures, the sub-panels respectively depict: (a) the skin conductance signal; (b) the deconvolved neural stimuli; (c) the arousal state estimate and its 95% confidence limits; (d) the probability of

impulse occurrence and its 95% confidence limits; (e) HAI. The background colors correspond to those in Fig. 10. On the Neurological Status Assessment Dataset, the two neural impulses at the beginning of the emotional stress period cause a considerable increase in HAI. HAI also remains high during the cognitive stressors and diminishes shortly thereafter during the relaxation period. The state estimates on the Driver Stress Dataset again yield no particular pattern depending on the road type and the increase in arousal during the second rest period is also visible again.

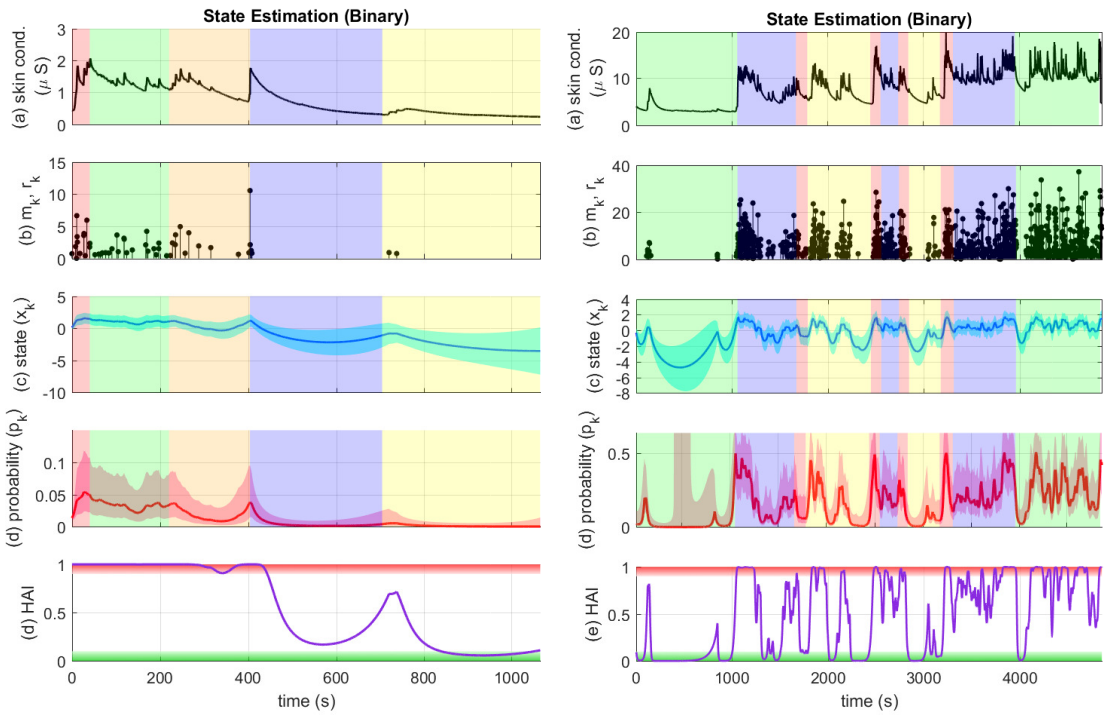

Figure 11: Sympathetic arousal estimation using model with one binary observation.

Results using the MPPSE method are shown in Fig. 12. Again, in both sub-figures, the sub-panels respectively depict: (a) the skin conductance signal; (b) the deconvolved neural stimuli; (c) the arousal state estimate and its 95% confidence limits; (d) the probability of impulse occurrence and its 95% confidence limits; (e) HAI. The background colors are also the same as in Fig. 10. Notably, on the Neurological Status Assessment Dataset, the

increase in HAI is not as prominent at the beginning of the emotional stress period. Results on the Driver Stress Dataset are similar to that of BOBSE. However, the MPP estimate is not as sensitive to an increase in HAI due to small neural impulses. In the region shortly before 1000 s, there are a few such small impulses. BOBSE yields a greater increase in HAI compared to the MPP estimate. Moreover, in the vicinity of 3000 s, the BOBSE method shows a moderate HAI increase followed by a second smaller HAI increase. However, on closer examination of the skin conductance signal and the detected impulses, the second increase is actually larger. The MPP estimate correctly places the larger HAI increase second.

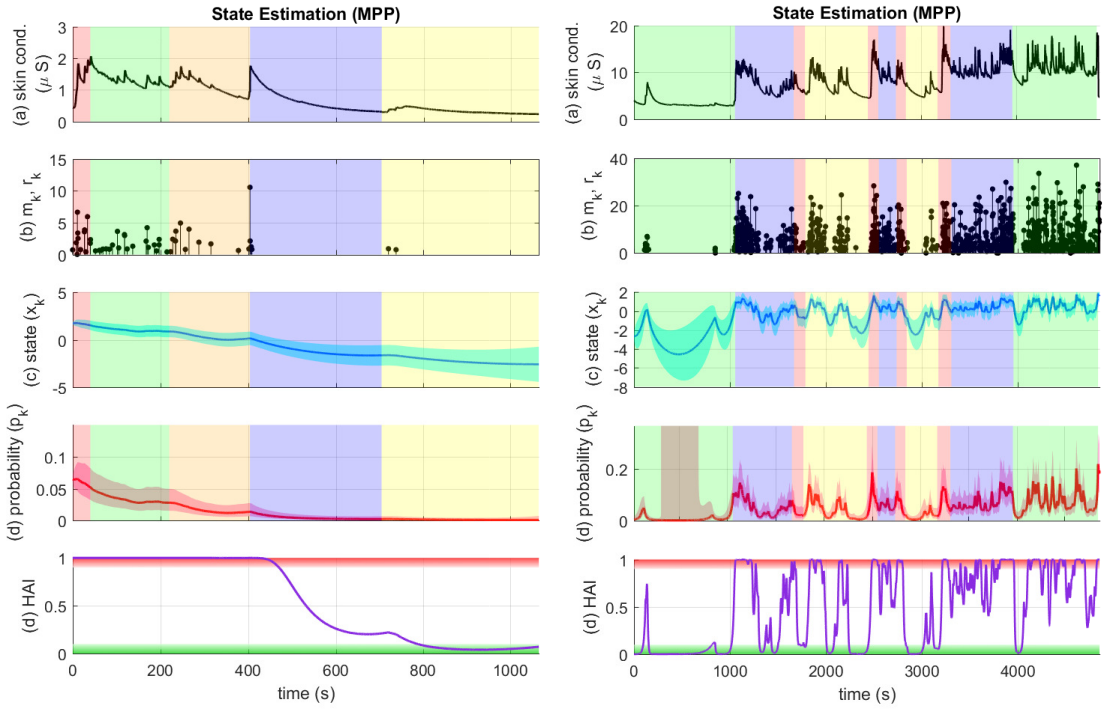

Figure 12: Sympathetic arousal estimation using model with one MPP observation.

SCR amplitudes contain sympathetic arousal information [83]. The BOBSE method does not incorporate the amplitude information of the SCRs (or equivalently of the underlying neural impulse bursts) for estimating sympathetic arousal from skin conductance.

While this method did provide good initial results, the lack of the amplitude information can result in arousal estimates that have exaggerated responses to small neural impulses. This is evident from the results for the BOBSE method in Fig. 11.

The BCOBSE method makes use of continuous-valued variables in addition to the binary SCR occurrences. However, there is a tendency to overfit since the location in the EM parameter space where the state fits almost perfectly to one of the continuous variables has a greater extrema. Therefore, the M-step updates related to the continuous variables need to be monitored and their updating halted when overfitting begins to occur. The method can also result in a very low σ_ε^2 estimate, a more jagged state estimate and sharp transitions in HAI. These effects are seen in Fig. 10.

Neural impulse amplitudes only exist at distinct points in time and not everywhere. Estimation based on treating the neural impulses as an MPP overcomes the sensitivity to small values and the tendency to overfit to a continuous variable. Consequently, only a small rise is seen at the beginning of the emotional stress period in the Neurological Status Assessment Dataset and a smoother state estimate compared to the BCOBSE method is seen in the Driver Stress Dataset. Thus the MPPSE method outperforms the BOBSE and BCOBSE methods in both these regards. Noise is a reality in any physiological signal. A noisy skin conductance signal can give rise to several small noisy impulses being detected during the deconvolution process and the MPP-based estimation scheme would be less sensitive to them in such scenarios.

2.5.3 Gaussian Approximation for Filter Derivation

We make a Gaussian approximation to the posterior density $p(x_k|y^k)$ in all of the filters we develop to calculate $x_{k|k}$ and $\sigma_{k|k}^2$. Therefore, it is necessary to first verify whether the EM algorithm is able to estimate x_k and recover model parameters accurately on simulated data (where the ground truth is known) before using it on experimental data. This is the procedure we follow for all of the major filters we develop. Based on results presented on

synthetic data in this chapter as well as in subsequent chapters, the filters based on making this Gaussian approximation to the posterior density do provide good estimation results.

One approach to improving the Gaussian approximation assumption is to use Gaussian kernels for the posterior density. One of the fundamental issues with our estimation problems, however, is that a ground truth is unavailable. Gaussian kernel methods for state-space filtering (e.g., [87, 88, 89, 90]) all require training data. In these methods, the posterior density is approximated as a weighted sum of Gaussian distributions via a kernelized representation. The parameters of these Gaussian distributions are learned through (x_k, y_k) pairs of training data. In our cases, we do not have any training data to begin with. Therefore, we are unable to use this Gaussian kernel method.

2.5.4 Expectation-maximization Algorithm Initialization

We randomly initialize the model parameters when running the EM algorithm. One of the model parameters that can be somewhat problematic is the process noise variance σ_ε^2. This term relates to how quickly the latent state we are tracking changes. Since sympathetic arousal is unobserved, it is somewhat difficult to know what σ_ε^2 is and how it may be best initialized. One approach to addressing this issue is to determine two *extremes* for σ_ε^2 and then check whether the random initialization and the extreme initializations result in convergence to the same model parameters.

With skin conductance data, we verified this using the Neurological Status Assessment Dataset and the MPPSE method. We first ran the EM algorithms during the cognitive stress periods and the relaxation periods separately for each subject. Thereafter, we took the two extreme values for σ_ε^2, re-initialized the EM algorithm at these values and ran it again. In general, the final parameters to which the algorithms converged were very close to what was obtained with the earlier random initialization. Moreover, the visual differences in the state estimates were also negligible.

We also performed the same verification for cortisol using experimental data from four

representative subjects using one of the Bayesian filters described in chapter 4. Here, the data was split into the sleep and awake periods. Again, however, the final estimates based on the two extreme value initializations of σ_ε^2 were very close to those obtained with the original random initialization and there were also only negligible visual differences in the state estimates.

2.6 Validation

We fundamentally address problems concerning the estimation of different unobserved state variables within the human body in this dissertation. As such, the ground truth is unknown, and this presents a number of difficulties. For instance, questions regarding how many features are to be selected for estimation, which features are optimal, which state-space model is the best and how the continuous-valued estimates are finally validated are all challenging to answer. Similar problems arise in biomedical applications elsewhere as well. For instance, mathematical modeling techniques that estimate central aortic blood pressure using only peripheral measurements have also encountered validation challenges [91]. In the case of skin conductance, our choice of features largely relies on the medical/psychophysiology literature. The rate of SCR occurrence, SCR amplitudes (which are related to the corresponding neural impulse occurrence rates and amplitudes) and the tonic levels are three of the most commonly used skin conductance indices of sympathetic arousal [70]. Hence, we chose to explore these features in our state-space models.

Validation/evaluation of the results relies on a general conformity to known physiology or a particular type of experiment (e.g., general agreement with stress conditions for skin conductance). One of our end objectives is to develop closed-loop controllers for patient care based on point process state-space models such as those described here. The performance of the closed-loop controllers could also be used to evaluate/validate the performance of the state-space models. An evaluation criteria could be established based on certain clinical endpoints and the best model selected from among them.

A possible option to validate/cross-check the results is to use another physiological signal(s) and examine its similarity with the state estimate. As a preliminary illustration of this possibility, we show the arousal estimation results for a participant who engaged in an n-back task (1-back and 3-back). An n-back task comprises of correctly recalling previous characters in a sequence of alphabetical letters displayed on a screen. The n-back tasks had to be performed under two conditions: (i) while listening to calming music; (ii) while listening to vexing music. The music was selected by the participant and there was a relaxation interval in-between the two sessions. We collected the participant's skin conductance from the non-dominant hand using Biopac sensors and functional Near Infrared Spectroscopy (fNIRS) using the NIRSport2 headset. The study protocol was approved by the University of Houston Institutional Review Board.

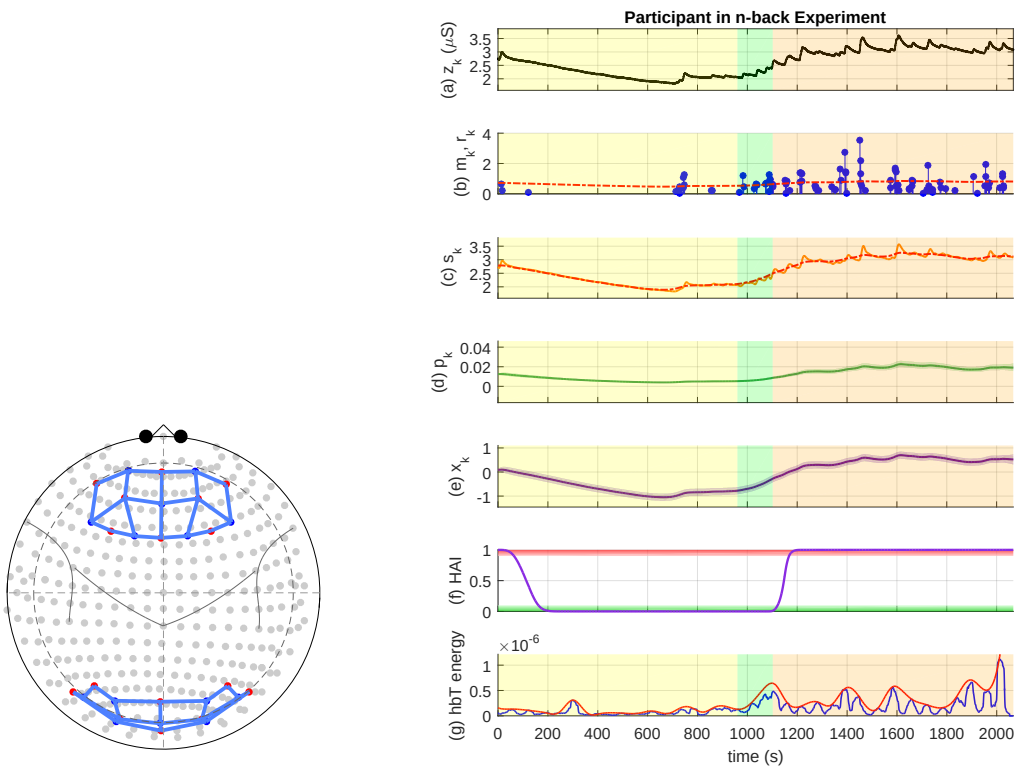

Figure 13: Sympathetic arousal estimation with fNIRS validation.

We processed the data using nirsLAB 2019.4. We first bandpass filtered the fNIRS signals between 0.01 and 0.2 Hz (default) and then computed the hemodynamic states. We next extracted the total hemoglobin (hbT) signals over the pre-frontal cortex (two bad channels were discarded based on a default nirsLAB coefficient of variation criterion of 7.5%) and calculated a sample-by-sample running average of the hbT squared energy over a 30 s window. Fig. 13 shows the skin conductance-based state estimates, and the mean and mean envelope of the hbT energies (based on all 20 channels over the pre-frontal cortex). In the figure, the sub-figure on the left depicts the fNIRS channel locations (the red and blue dots denote the sources and detectors respectively). In the sub-figure on the right, the sub-panels respectively depict: (a) the skin conductance signal z_k; (b) the deconvolved neural impulses (blue) and the estimated fit to the amplitudes (red); (c) the tonic component (orange) and its estimated fit (red); (d) the probability of impulse occurrence and its 95% confidence limits; (e) the arousal state estimate and its 95% confidence limits; (f) HAI calculated based on the median state (the regions above 90% and below 10% are shaded in red and green respectively); (g) the mean (blue) and mean envelope (red) of the hbT energies. The background colors correspond to performing the n-back trials during calming music (yellow), relaxation (green) and performing the n-back trials during vexing music (orange). Here we used a state-space model with one MPP and one continuous-valued variable for estimation that is introduced in chapter 4 (the tonic component was the continuous-valued variable). A similarity is seen between the skin conductance-based sympathetic arousal estimates and the hbT mean energies. Both x_k and the hbT energies are lower during calming music and are higher during vexing music. It is likely that the participant had to put in more cognitive effort during the vexing music session to get the answers correct. Since a ground truth is unavailable for the estimation problems we describe here, the use of other physiological signals such as fNIRS for validation is an option. Moreover, since the size of the pupil is also influenced by the sympathetic nervous system, validating the estimate against a pupil size signal is also a viable alternative.

3 Sympathetic Arousal Estimation using Skin Conductance and Heart Rate Data

3.1 Modeling and Autonomic Regulation of Heart Rate

We next focus on extracting sympathetic arousal information from the heart to supplement the skin conductance features considered earlier. Recall that nerve fibers from the sympathetic branch of the autonomic nervous system innervate the sweat glands. Consequently, skin conductance becomes a sensitive index of sympathetic arousal [61]. It is more complicated with the heart. The heart is innervated by both the sympathetic and parasympathetic branches which have a "push-pull" effect on it. Moreover, the heart also has its *own* pacing mechanism [92]. Now an EKG is one of the most common methods of extracting information regarding cardiac activity. An EKG signal comprises of repeating elements known as QRS complexes. Most prominent in a QRS complex is the R-peak which accompanies ventricular contraction. Heart rate can be easily obtained from an EKG signal by extracting the RR-intervals and calculating the time difference between successive beats. In this chapter, we describe two different state-space models for estimating sympathetic arousal based on skin conductance and heart rate measurements. The first, described in [93], utilizes binary SCR occurrences and RR-intervals, and the second, described in [94], utilizes binary SCR occurrences, tonic and phasic information (as continuous-valued observations) and RR-intervals.

The sympathetic nervous system increases heart rate and the force of contraction via the neurotransmitter norepinephrine [95]. In contrast, parasympathetic activation causes the release of actylcholine at the heart and has the opposite effect. Beat-to-beat variations in RR-intervals, known as heart rate variability (HRV), reflect changes in sympathetic and parasympathetic control on the heart. Here we relate sympathetic arousal to heart rate. Studies in animal models have shown that the stimulation of autonomic nerve fibers leading

to the heart results in an almost linear relationship between stimulation frequency and RR-intervals [96, 97]. Based on findings in these studies, we select a linear model to capture the relationship between RR-intervals and sympathetic arousal x_k. Note, however, that we do not assume that RR-intervals are linear, but rather that the *effect* of sympathetic arousal *on* RR-intervals is linear. It is, however, possible to assume more complex effects between x_k and RR-intervals if necessary.

Heartbeats occur due to the depolarization of cells in the heart's sinoatrial (SA) node which subsequently propagates throughout the atria and ventricles. The rise in the membrane potentials in the SA node cells can be modeled as a Gaussian random walk with drift [98, 99]. Consequently, the times between successive ventricular contractions can be modeled using the inverse Gaussian probability density model. Barbieri et al., [98, 100] successfully used a history-dependent inverse Gaussian (HDIG) probability density function to model RR-intervals. If L consecutive R-peaks occur at times u_l within $(0, T]$ such that $0 < u_1 < u_2 < \ldots < u_L \leq T$, and $h_l = u_l - u_{l-1}$ is the l^{th} RR-interval, then the HDIG density function for RR-intervals at $t > u_l$ is,

$$g(t|u_l) = \sqrt{\frac{\theta_{q+1}}{2\pi(t-u_l)^3}}\ \exp\left\{\frac{-\theta_{q+1}[t-u_l-\hat{\mu}]^2}{2\hat{\mu}^2(t-u_l)}\right\}, \tag{40}$$

where q is the model order, θ_{q+1} is related to the variance and

$$\hat{\mu} = \theta_0 + \sum_{i=1}^{q} \theta_i h_{l-i+1} \tag{41}$$

is the HDIG mean [100]. The θ_i's are coefficients to be determined. This model expresses the dependence of an RR-interval on its immediate history (this dependence has also led to the successful application of autoregressive models to the analysis of HRV [101, 102]). Barbieri et al., [98] divided the time axis into bins of size $\Delta = 5$ ms and performed local likelihood estimation to determine the θ_i's every Δ ms (i.e., the θ_i's were time-varying). These time-varying parameters capture part of the non-stationary nature of HRV that occurs due to

underlying pathological and physiological reasons [103].

Based on our earlier assumption of linearity between sympathetic arousal x_k and the RR-intervals, we re-define a new HDIG RR-interval mean μ as one which depends linearly on both the immediate history and x_k as

$$\mu = \theta_0 + \sum_{i=1}^{q} \theta_i h_{l-i+1} + \eta x_k, \tag{42}$$

where η is a coefficient to be determined. Moreover, we also assume that the θ_i's are fixed and that variations in sympathetic arousal account for part of the RR-interval stochasticity. According to this formulation, changes in arousal would cause the HDIG probability density function to shift to the right or the left.

We typically analyze skin conductance data at a sampling frequency of 4 Hz due to its low bandwidth (sampling time $t_s = 250$ ms). Now the bin size $\Delta = 5$ ms proposed by Barbieri et al., [98] used for the HDIG model is much smaller. There are $J = t_s/\Delta = 50$ heart rate observation bins corresponding to the k^{th} skin conductance sample. We index these smaller heart rate bins over j and generate a binary point process by assigning $n_{k,j} = 1$ or $n_{k,j} = 0$ depending on whether or not an R-peak occurred at the time. The joint density over these J observations is then [104]

$$P(n_{k,1}, n_{k,2}, \ldots, n_{k,J}|x_k) = e^{\sum_{j=1}^{J} \log(\lambda_{k,j}\Delta) n_{k,j} - \lambda_{k,j}\Delta}, \tag{43}$$

where the conditional intensity function (CIF) $\lambda_{k,j}$ is

$$\lambda_{k,j} \triangleq \frac{g(t_{k,j}|u_{k,j})}{1 - \int_{u_{k,j}}^{t_{k,j}} g(z|u_{k,j})dz}, \tag{44}$$

and $u_{k,j}$ is the time of occurrence of the last R-peak prior to $t_{k,j}$. The sequence

$$\mathcal{N}^K = \{n_{1,1}, n_{1,2}, \ldots, n_{1,J}, n_{2,1}, n_{2,2}, \ldots, n_{2,J}, \ldots, n_{K,1}, n_{K,2}, \ldots, n_{K,J}\} \tag{45}$$

forms the set of heart rate observations. We use $\mathcal{N}^K$ as an additional feature capturing sympathetic arousal information in the state-space models described here.

3.2 Data

3.2.1 Smart Reasoning for Well-being at Home and at Work Dataset

We first use the Smart Reasoning for Well-being at Home and at Work (SWELL) Dataset [105]. In the SWELL experimental study, skin conductance, EKG, computer interactions, video and posture data were recorded from subjects as they engaged in activities typically performed in an office. The tasks included preparing reports and presentations, and responding to e-mails. The experimental trials were divided into three blocks and each subject had to prepare two reports and one Powerpoint presentation (on one of those reports) in each block. In the first block, subjects could take as much time as they liked. In one of the other two blocks, subjects were sent eight e-mails (some of which were irrelevant) and were instructed to make use of the incoming information and reply as required. In the remaining block, a subject could only take $2/3^{\text{rds}}$ of the time taken for the very first block. The order of these last two blocks was varied from subject to subject. Here, we label these blocks as the Unfamiliar Task (UT) block (since the subjects were newly introduced to the experimental task), the Repeated Task (RT) block and the Constrained Task (CT) block.

The six report and presentation topics included: stress at work, healthy living, internet privacy, tourist attractions in Perth, the life of Napoleon and details for a road trip across the United States. Two topics were randomly selected for each block and the subjects were permitted to browse the web in preparing their reports. Each block was preceded by 8 min of relaxing music. Here, we analyze the data from one subject (subject 16) over all three blocks.

The skin conductance and EKG data were recorded at 2048 Hz in the SWELL dataset. Many of the skin conductance recordings were heavily noise corrupted. Subject 16 was selected due to the least signal contamination being present across blocks UT, RT and CT.

We first manually identified locations where motion artifact corruption was present (usually manifesting as sharp drops) and linearly interpolated over them. We finally lowpass filtered the skin conductance at 0.5 Hz similar to [106] and downsampled the data to 4 Hz. We also performed peak detection on the EKGs to identify R-peaks.

3.2.2 Fear Conditioning Dataset

We secondly use data from a Pavlovian fear conditioning experiment. In this general category of experiments, a subject learns an association between a biologically-relevant stimulus and a neutral cue through repeated pairing [107]. The paradigm originated with experiments conducted by the Russian scientist Ivan Pavlov in the early part of the 20^{th} century [107]. In his classic experiment, Pavlov repeatedly paired the ringing of a bell with food, eventually causing a dog to salivate merely at the ringing of the bell [108]. The food was named as the unconditioned stimulus (US) and the ringing of the bell as the conditioned stimulus (CS). Learning the association between the CS and US lies at the heart of Pavlovian conditioning; eventually, the CS alone will begin to elicit the biological response typically associated with the US. In fear conditioning experiments, the US is unpleasant. It can take the form of a mild electric shock, a loud sound, an aversive image, a blast of air to the throat etc. [108, 109]. Pavlovian fear conditioning has been examined in both human subjects and animal models. Additional forms of fear conditioning experiments arose later. These include differential conditioning and the use of more complex stimuli. In differential conditioning, there are two types of conditioned stimuli—CS+ and CS-. The CS- is never associated with the US. The CS+ may be associated fully or partially with the US. The CS+ can also be chosen to reinforce the threat of the US (e.g., the image of a fearful face may be used as the CS+ and a neutral face as the CS-).

Here we use the PsPM-TC: SCR, EKG, EMG and Respiration Measurements in a Discriminant Trace Fear Conditioning Task with Visual CS and Electrical US Dataset [110]. The dataset is described in detail in [111, 112, 113], and is publicly available through the

Zenodo repository. The original experiment involved 23 subjects (13 males, 10 females, age 23.8 ± 3.0 years) from whom four subjects were discarded [111]. The dataset available online contains physiological signal measurements from the other 19 healthy subjects. In trace fear conditioning, as opposed to delay fear conditioning, there is a time gap between the termination of the CS+ cue and the US onset [108]. Blue and red rectangles on a computer screen were used as the CS+ and CS- visual cues. The US was a series of 0.2 ms square electrical pulses applied at a frequency of 10 Hz for a duration of 0.5 s to the subject's forearm using a pin-cathode/ring-anode electrode configuration. The stimulation intensity was set to approximately 90% of each subject's pain threshold following a two-step procedure. Skin conductance was recorded from Ag/AgCl cup electrodes placed on the thenar/hypothenar of each subject's non-dominant hand and EKG was likewise recorded using Ag/AgCl electrodes placed on the limbs. Only 50% of the CS+ trials were accompanied by the US. Skin conductance and EKG can be contaminated by various sources of noise including motion artifacts and powerline noise. We analyzed data from 12 subjects for whom only low to moderate noise contamination was present in the signals.

3.3 Methods

As stated earlier, we developed two different state-space models and corresponding Bayesian filters to estimate sympathetic arousal from skin conductance and heart rate. The model that only utilizes binary SCR occurrences and RR-intervals [93] can be considered to be a special case of the extended model [94]. The filter equations for the simpler model can also be derived by taking the extended model and dropping the terms corresponding to r_k and s_k in the posterior density $p(x_k|y^k)$. We evaluated the extended model on the PsPM-TC Dataset and the simpler one on the SWELL Dataset. Note also that when heart rate observations are included, a slight modification is necessary to the EM algorithm. Owing to computational complexity, we are unable to estimate η at the M-step and instead have to rely on an alternate means of estimation. The EM algorithm details are described

below.

3.3.1 State-space Model with One Binary, Two Continuous and One Spiking-type Observation

Here we assume that sympathetic arousal varies with time as

$$x_k = \rho x_{k-1} + \alpha I_k + \varepsilon_k \tag{46}$$

similar to eq. (1). We set $I_k = 1$ at the times corresponding to the presentation of the CS+, CS- and US stimuli. We also assume that x_k is related to the probability of SCR occurrence p_k through

$$p_k = \frac{1}{1 + e^{-(\beta_0 + \beta_1 x_k)}}, \tag{47}$$

where β_0 and β_1 are coefficients to be determined.

The derivation of the E-step equations is similar to what was seen in the earlier chapter. We make a Gaussian approximation to the posterior density $p(x_k|y^k)$ similar to [46] to obtain the following filter equations:

$$x_{k|k-1} = \rho x_{k-1|k-1} + \alpha I_k, \tag{48}$$

$$\sigma^2_{k|k-1} = \rho^2 \sigma^2_{k-1|k-1} + \sigma^2_\varepsilon, \tag{49}$$

$$C_k = \frac{\sigma^2_{k|k-1}}{\sigma^2_v \sigma^2_w + \sigma^2_{k|k-1}(\gamma_1^2 \sigma^2_w + \delta_1^2 \sigma^2_v)}, \tag{50}$$

$$\begin{aligned} x_{k|k} = x_{k|k-1} + C_k \Bigg[& \beta_1 \sigma^2_v \sigma^2_w (m_k - p_{k|k}) + \gamma_1 \sigma^2_w (r_k - \gamma_0 - \gamma_1 x_{k|k-1}) \\ & + \delta_1 \sigma^2_v (s_k - \delta_0 - \delta_1 x_{k|k-1}) + \sigma^2_v \sigma^2_w \sum_{j=1}^{J} \frac{1}{\lambda_{k,j|k}} \frac{\partial \lambda_{k,j|k}}{\partial x_k} (n_{k,j} - \lambda_{k,j|k} \Delta) \Bigg], \end{aligned} \tag{51}$$

and

$$\sigma^2_{k|k} = \left\{ \frac{1}{\sigma^2_{k|k-1}} + \beta_1^2 p_{k|k}(1-p_{k|k}) + \frac{\gamma_1^2}{\sigma_v^2} + \frac{\delta_1^2}{\sigma_w^2} - \sum_{j=1}^{J} \left[\frac{1}{\lambda_{k,j|k}} \frac{\partial^2 \lambda_{k,j|k}}{\partial x_k^2} (n_{k,j} - \lambda_{k,j|k}\Delta) - \frac{n_{k,j}}{\lambda^2_{k,j|k}} \left(\frac{\partial \lambda_{k,j|k}}{\partial x_k} \right)^2 \right] \right\}^{-1}. \tag{52}$$

The M-step is more complicated. Ideally, all model parameters related to both skin conductance and heart rate should be estimated at the M-step simultaneously. Recall that we have to determine $\theta_0, \theta_1, \ldots, \theta_{q+1}$ and η for heart rate. Calculating these values at the M-step requires the maximization of [46]

$$\bar{Q} = \sum_{k=1}^{K} \sum_{j=1}^{J} \mathbb{E}\big[\log(\lambda_{k,j}\Delta) n_{k,j} - \lambda_{k,j}\Delta \big]. \tag{53}$$

Maximizing $\bar{Q}$ with respect to the θ_i's and η for a fixed model order q is extremely time consuming. Additionally, multiple values of q need to be evaluated for selecting the best order. Owing to this large time complexity, we resorted to an alternate two-step strategy for determining the model parameters related to heart rate.

- Step 1: Determining the model order q and the θ_i coefficients

  The HDIG density function models the RR-interval mean as a weighted sum of the previous q RR-intervals. This is similar to an autoregressive model where a value in a time series is predicted based on its past values [103]. Barbieri et al., [98, 100] assessed goodness-of-fit when modeling RR-intervals based on the HDIG density function using the Kolmogorov-Smirnov (KS) plot. The KS plot is based on the time rescaling theorem [114] and provides an indication of how well a CIF fits to point process observations. In addition to heartbeat data, the time rescaling theorem is also frequently used in the analysis of neural spike trains [115, 116, 117]. The closer the KS plot is to the 45° diagonal, the better the fit is to the point process observations. Thus,

the maximum distance between the KS plot and the 45° diagonal (known as the KS distance), provides a quantitative measure of goodness-of-fit. For each model order q, we estimated $\theta_0, \theta_1, \ldots, \theta_{q+1}$ offline via maximum likelihood [98]. The best model order q and the θ_i coefficients were selected based on the smallest KS distance. We performed step 1 for each participant prior to arousal estimation using EM.

- Step 2: Determining η

  After selecting the θ_i's and the model order q, we run the full EM algorithm for arousal estimation for a fixed set of values for η. Since RR-intervals decrease (i.e., heart rate speeds up) with increased sympathetic drive, we chose to try a set of negative values for η. We resorted to this two-step strategy because determining η and the θ_i's simultaneously at the M-step proved to be extremely time consuming. $\bar{Q}$ in (53) can be approximated by

$$\bar{Q} \approx \sum_{k=1}^{K}\sum_{j=1}^{J} \log(\lambda_{k,j|K}\Delta)n_{k,j} - \lambda_{k,j|K}\Delta + \frac{1}{2}\left[\frac{1}{\lambda_{k,j|K}}\frac{\partial^2\lambda_{k,j|K}}{\partial x_k^2}(n_{k,j} - \lambda_{k,j|K}\Delta) - \frac{n_{k,j}}{\lambda_{k,j|K}^2}\left(\frac{\partial\lambda_{k,j|K}}{\partial x_k}\right)^2\right]\sigma_{k|K}^2. \tag{54}$$

  Ideally, the η value with the largest $\bar{Q}$ during state estimation should be selected. However, the inclusion of x_k in the HDIG model causes the KS plot obtained via maximum likelihood to change. Therefore, η should be chosen to maximize $\bar{Q}$ subject to the new KS plot falling within or reasonably close to the 95% confidence limits.

In summary, the EM algorithm is split into two. All model parameters excluding the ones related to heart rate are calculated in the regular EM manner. The θ_i coefficients are estimated offline prior to running the EM algorithm. For η, we run the EM algorithm for different values of η and then select the one with the highest log-likelihood subject to the KS plot falling reasonably close to the 45° diagonal.

3.3.2 State-space Model with One Binary and One Spiking-type Observation

Since no external inputs such as such as cues to electric shocks are present in the SWELL experimental study, we choose the random walk $x_k = x_{k-1} + \varepsilon_k$ for the state equation. We also consider x_k to be related to p_k through

$$p_k = \frac{1}{1 + e^{-(\beta_0 + x_k)}}, \tag{55}$$

where β_0 is determined empirically.

The filter equations are as follows:

$$x_{k|k-1} = x_{k-1|k-1}, \tag{56}$$

$$\sigma^2_{k|k-1} = \sigma^2_{k-1|k-1} + \sigma^2_{\varepsilon}, \tag{57}$$

$$x_{k|k} = x_{k|k-1} + \sigma^2_{k|k-1}\left[(m_k - p_{k|k}) + \sum_{j=1}^{J} \frac{1}{\lambda_{k,j|k}} \frac{\partial \lambda_{k,j|k}}{\partial x_k}(n_{k,j} - \lambda_{k,j|k}\Delta)\right], \tag{58}$$

and

$$\sigma^2_{k|k} = \left\{ \frac{1}{\sigma^2_{k|k-1}} + p_{k|k}(1 - p_{k|k}) \right. \\ \left. - \sum_{j=1}^{J} \left[\frac{1}{\lambda_{k,j|k}} \frac{\partial^2 \lambda_{k,j|k}}{\partial x_k^2}(n_{k,j} - \lambda_{k,j|k}\Delta) - \frac{n_{k,j}}{\lambda^2_{k,j|k}} \left(\frac{\partial \lambda_{k,j|k}}{\partial x_k} \right)^2 \right] \right\}^{-1}. \tag{59}$$

The equations are similar to those in the previous case except that now the terms for r_k and s_k are absent.

3.4 Derivation of Expectation-maximization Algorithm Equations for Model with One Binary, Two Continuous and One Spiking-type Observation

At this point we briefly digress from our usual flow to provide the EM algorithm derivations for the state-space model with one binary, two continuous and one spiking-type observation. Mathematically, this represents the most complicated case without an MPP observation. All other filter equations can be derived by making adjustments to this particular state-space model. This section also forms a self-contained unit having just the EM derivations.

3.4.1 E-step Filter Equations

We follow an approach similar to [46] in deriving the filter updates. Recall that we have the following linear relationships between x_k, r_k and s_k:

$$r_k = \gamma_0 + \gamma_1 x_k + v_k \tag{60}$$

and

$$s_k = \delta_0 + \delta_1 x_k + w_k. \tag{61}$$

We take the two noise terms v_k and w_k to be independent of each other. Consequently, the density functions $p(r_k|x_k)$ and $p(s_k|x_k)$ conditioned on already having observed x_k are also independent of each other. Taking y^k to denote the observations recorded up to time index k, we have

$$\begin{aligned} p(x_k|y^k) &= \frac{p(x_k|y^{k-1})p(y_k|x_k)}{p(y_k|y^{k-1})} \\ &= \frac{p(x_k|y^{k-1})P(m_k|x_k)p(r_k|x_k)p(s_k|x_k)P(n_{k,1:J}|x_k)}{p(y_k|y^{k-1})} \end{aligned}$$

$$\propto \exp\Bigg[\frac{-(x_k - x_{k|k-1})^2}{2\sigma^2_{k|k-1}} + m_k \log\left(\frac{p_k}{1-p_k}\right) + \log(1-p_k)$$
$$- \frac{(r_k - \gamma_0 - \gamma_1 x_k)^2}{2\sigma_v^2} - \frac{(s_k - \delta_0 - \delta_1 x_k)^2}{2\sigma_w^2} + \sum_{j=1}^{J} \log(\lambda_{k,j}\Delta) n_{k,j} - \lambda_{k,j}\Delta\Bigg]. \quad (62)$$

Therefore,

$$\log\Big[p(x_k|y^k)\Big] = \Bigg[\frac{-(x_k - x_{k|k-1})^2}{2\sigma^2_{k|k-1}} + m_k \log\left(\frac{p_k}{1-p_k}\right) + \log(1-p_k)$$
$$- \frac{(r_k - \gamma_0 - \gamma_1 x_k)^2}{2\sigma_v^2} - \frac{(s_k - \delta_0 - \delta_1 x_k)^2}{2\sigma_w^2} + \sum_{j=1}^{J} \log(\lambda_{k,j}\Delta) n_{k,j} - \lambda_{k,j}\Delta\Bigg]$$
$$+ \text{const.} \quad (63)$$

We make a Gaussian approximation to the posterior $p(x_k|y^k)$ to derive the mean and variance. Setting the partial derivative of the logarithm term to 0, we obtain

$$\frac{\partial}{\partial x_k}\log\Big[p(x_k|y^k)\Big] = \frac{-(x_k - x_{k|k-1})}{\sigma^2_{k|k-1}} + \beta_1(m_k - p_k) + \frac{\gamma_1(r_k - \gamma_0 - \gamma_1 x_k)}{\sigma_v^2}$$
$$+ \frac{\delta_1(s_k - \delta_0 - \delta_1 x_k)}{\sigma_w^2} + \sum_{j=1}^{J} \frac{1}{\lambda_{k,j|k}} \frac{\partial \lambda_{k,j|k}}{\partial x_k}(n_{k,j} - \lambda_{k,j|k}\Delta) = 0. \quad (64)$$

Solving for x_k in the equation above provides the filter update (i.e., the mean value) for $x_{k|k}$. Here we have taken

$$\frac{\partial p_k}{\partial x_k} = \beta_1 p_k(1-p_k) \quad (65)$$

when calculating the partial derivative. Similarly, the second partial derivative is

$$\frac{\partial^2}{\partial x_k^2}\log\Big[p(x_k|y^k)\Big] = \frac{-1}{\sigma^2_{k|k-1}} - \beta_1 \frac{\partial p_k}{\partial x_k} - \frac{\gamma_1^2}{\sigma_v^2} - \frac{\delta_1^2}{\sigma_w^2}$$
$$+ \frac{\partial}{\partial x_k}\left[\sum_{j=1}^{J} \frac{1}{\lambda_{k,j|k}} \frac{\partial \lambda_{k,j|k}}{\partial x_k}(n_{k,j} - \lambda_{k,j|k}\Delta)\right]$$

$$= \frac{-1}{\sigma^2_{k|k-1}} - \beta_1^2 p_k(1-p_k) - \frac{\gamma_1^2}{\sigma_v^2} - \frac{\delta_1^2}{\sigma_w^2} + \sum_{j=1}^{J}\left[\frac{1}{\lambda_{k,j|k}}\frac{\partial^2\lambda_{k,j|k}}{\partial x_k^2}(n_{k,j}-\lambda_{k,j|k}\Delta) - \frac{n_{k,j}}{\lambda^2_{k,j|k}}\left(\frac{\partial\lambda_{k,j|k}}{\partial x_k}\right)^2\right]. \quad (66)$$

The filter update for $\sigma^2_{k|k}$ is given by [46]

$$\sigma^2_{k|k} = \left\{ -\frac{\partial^2}{\partial x_k^2}\log\Big[p(x_k|y^k)\Big]\right\}^{-1}. \quad (67)$$

3.4.2 M-step Updates

Complete Data Log-likelihood

Taking Θ to denote all the model parameters, the complete data likelihood conditioned on Θ is

$$\begin{aligned} p(\mathcal{Y}^K, \mathcal{X}^K|\Theta) = &\prod_{k=1}^{K} p_k^{m_k}(1-p_k)^{1-m_k} \times \prod_{k=1}^{K}\frac{1}{\sqrt{2\pi\sigma_v^2}}e^{-\frac{(r_k-\gamma_0-\gamma_1 x_k)^2}{2\sigma_v^2}} \\ &\times \prod_{k=1}^{K}\frac{1}{\sqrt{2\pi\sigma_w^2}}e^{-\frac{(s_k-\delta_0-\delta_1 x_k)^2}{2\sigma_w^2}} \times \prod_{k=1}^{K} e^{\sum_{j=1}^{J}\log(\lambda_{k,j}\Delta)n_{k,j}-\lambda_{k,j}\Delta} \\ &\times \prod_{k=1}^{K}\frac{1}{\sqrt{2\pi\sigma_\varepsilon^2}}e^{-\frac{(x_k-\rho x_{k-1}-\alpha I_k)^2}{2\sigma_\varepsilon^2}}. \end{aligned} \quad (68)$$

Therefore, the expected log-likelihood is

$$\begin{aligned} Q = &\sum_{k=1}^{K}\mathbb{E}\Big[m_k(\beta_0+\beta_1 x_k) - \log\left(1+e^{\beta_0+\beta_1 x_k}\right)\Big] + \frac{(-K)}{2}\log\left(2\pi\sigma_v^2\right) \\ &- \sum_{k=1}^{K}\frac{\mathbb{E}\Big[(r_k-\gamma_0-\gamma_1 x_k)^2\Big]}{2\sigma_v^2} + \frac{(-K)}{2}\log\left(2\pi\sigma_w^2\right) - \sum_{k=1}^{K}\frac{\mathbb{E}\Big[(s_k-\delta_0-\delta_1 x_k)^2\Big]}{2\sigma_w^2} \\ &+ \sum_{k=1}^{K}\sum_{j=1}^{J}\mathbb{E}\Big[\log(\lambda_{k,j}\Delta)n_{k,j} - \lambda_{k,j}\Delta\Big] + \frac{(-K)}{2}\log\left(2\pi\sigma_\varepsilon^2\right) \\ &- \sum_{k=1}^{K}\frac{\mathbb{E}\Big[(x_k-\rho x_{k-1}-\alpha I_k)^2\Big]}{2\sigma_\varepsilon^2}. \end{aligned} \quad (69)$$

Similar to [41], we denote the following expected values as:

$$x_{k|K} = \mathbb{E}\Big[x_k|y^K, \Theta\Big], \tag{70}$$

$$U_k = \mathbb{E}\Big[x_k^2|y^K, \Theta\Big], \tag{71}$$

and

$$U_{k,k+1} = \mathbb{E}\Big[x_k x_{k+1}|y^K, \Theta\Big]. \tag{72}$$

M-step Updates for α and ρ

Let

$$Q_1 = \frac{1}{2\sigma_\varepsilon^2}\sum_{k=1}^{K}\mathbb{E}\Big[(x_k - \rho x_{k-1} - \alpha I_k)^2\Big] \tag{73}$$

denote the term in Q containing α and ρ. While it is possible to determine the starting state x_0 as a separate parameter, we follow one of the options in [41, 44] and set $x_0 = x_1$. This permits some bias at the beginning. Therefore,

$$Q_1 = \frac{1}{2\sigma_\varepsilon^2}\left\{\sum_{k=2}^{K}\mathbb{E}\Big[(x_k - \rho x_{k-1} - \alpha I_k)^2\Big] + \mathbb{E}\Big[(\alpha I_1)^2\Big]\right\}. \tag{74}$$

We take the partial derivatives of Q_1 with respect to α and ρ and set them to 0 to obtain the M-step updates. Taking the partial derivative of Q_1 with respect to α, we have

$$\frac{\partial Q_1}{\partial \alpha} = \frac{1}{2\sigma_\varepsilon^2}\left\{\sum_{k=2}^{K}\mathbb{E}\big[-2I_k(x_k - \rho x_{k-1} - \alpha I_k)\big] + 2\alpha I_1^2\right\}. \tag{75}$$

Setting this term to 0 yields

$$0 = -\sum_{k=2}^{K} I_k\mathbb{E}\big[x_k\big] + \rho\sum_{k=2}^{K} I_k\mathbb{E}\big[x_{k-1}\big] + \alpha\sum_{k=1}^{K} I_k^2$$

$$= -\sum_{k=2}^{K} I_k x_{k|K} + \rho \sum_{k=2}^{K} I_k x_{k-1|K} + \alpha \sum_{k=1}^{K} I_k^2. \tag{76}$$

Taking the partial derivative of Q_1 with respect to ρ yields

$$\frac{\partial Q_1}{\partial \rho} = \frac{1}{2\sigma_\varepsilon^2} \left\{ \sum_{k=2}^{K} \mathbb{E}\big[-2x_{k-1}(x_k - \rho x_{k-1} - \alpha I_k) \big] \right\}. \tag{77}$$

Setting this term to 0, we obtain

$$\begin{aligned} 0 &= -\sum_{k=2}^{K} \mathbb{E}\big[x_k x_{k-1}\big] + \rho \sum_{k=2}^{K} \mathbb{E}\big[x_{k-1}^2\big] + \alpha \sum_{k=2}^{K} I_k \mathbb{E}\big[x_{k-1}\big] \\ &= -\sum_{k=1}^{K-1} U_{k,k+1} + \rho \sum_{k=1}^{K-1} U_k + \alpha \sum_{k=2}^{K} I_k x_{k-1|K}. \end{aligned} \tag{78}$$

The solutions to these two simultaneous equations provide α and ρ.

M-step Updates for γ_0, γ_1, δ_0 and δ_1

Let

$$Q_2 = \sum_{k=1}^{K} \frac{\mathbb{E}\Big[(r_k - \gamma_0 - \gamma_1 x_k)^2\Big]}{2\sigma_v^2} \tag{79}$$

denote the term in Q containing γ_0 and γ_1. Taking the partial derivative with respect to γ_0 we obtain

$$\frac{\partial Q_2}{\partial \gamma_0} = \frac{1}{2\sigma_v^2} \sum_{k=1}^{K} -2\mathbb{E}\big[r_k - \gamma_0 - \gamma_1 x_k\big]. \tag{80}$$

Setting this term to 0 yields

$$0 = -\sum_{k=1}^{K} r_k + \gamma_0 K + \gamma_1 \sum_{k=1}^{K} \mathbb{E}\big[x_k\big] = -\sum_{k=1}^{K} r_k + \gamma_0 K + \gamma_1 \sum_{k=1}^{K} x_{k|K}. \tag{81}$$

Taking the partial derivative of Q_2 with respect to γ_1 yields

$$\frac{\partial Q_2}{\partial \gamma_1} = \frac{1}{2\sigma_v^2} \sum_{k=1}^{K} -2\mathbb{E}\big[x_k(r_k - \gamma_0 - \gamma_1 x_k)\big]. \tag{82}$$

Setting this term to 0, we obtain

$$\begin{aligned} 0 &= -\sum_{k=1}^{K} r_k \mathbb{E}\big[x_k\big] + \gamma_0 \sum_{k=1}^{K} \mathbb{E}\big[x_k\big] + \gamma_1 \sum_{k=1}^{K} \mathbb{E}\big[x_k^2\big] \\ &= -\sum_{k=1}^{K} r_k x_{k|K} + \gamma_0 \sum_{k=1}^{K} x_{k|K} + \gamma_1 \sum_{k=1}^{K} U_k. \end{aligned} \tag{83}$$

The solutions to these two simultaneous equations provide γ_0 and γ_1. δ_0 and δ_1 may be obtained similarly from the term in Q containing s_k. In the case of an MPP observation, the constant coefficients are calculated by performing the corresponding summations only over the reduced set of indices at which $m_k = 1$.

M-step Updates for σ_v^2 and σ_w^2

Let

$$Q_3 = \frac{-K}{2} \log\left(2\pi\sigma_v^2\right) - \sum_{k=1}^{K} \frac{\mathbb{E}\Big[(r_k - \gamma_0 - \gamma_1 x_k)^2\Big]}{2\sigma_v^2} \tag{84}$$

denote the term in Q containing σ_v^2. Setting the partial derivative of Q_3 with respect to σ_v^2 to 0, we obtain

$$\frac{\partial Q_3}{\partial \sigma_v^2} = \frac{-K}{2\sigma_v^2} + \frac{1}{2\sigma_v^4} \sum_{k=1}^{K} \mathbb{E}\Big[(r_k - \gamma_0 - \gamma_1 x_k)^2\Big] = 0. \tag{85}$$

Therefore,

$$\begin{aligned} \sigma_v^2 &= \frac{1}{K} \sum_{k=1}^{K} \mathbb{E}\Big[(r_k - \gamma_0 - \gamma_1 x_k)^2\Big] \\ &= \frac{1}{K} \Bigg\{ \sum_{k=1}^{K} r_k^2 + K\gamma_0^2 + \gamma_1^2 \sum_{k=1}^{K} \mathbb{E}\big[x_k^2\big] - 2\gamma_0 \sum_{k=1}^{K} r_k - 2\gamma_1 \sum_{k=1}^{K} r_k \mathbb{E}\big[x_k\big] \end{aligned}$$

$$+ 2\gamma_0\gamma_1 \sum_{k=1}^{K} \mathbb{E}\big[x_k\big] \Bigg\}$$
$$= \frac{1}{K}\Bigg\{ \sum_{k=1}^{K} r_k^2 + K\gamma_0^2 + \gamma_1^2 \sum_{k=1}^{K} U_k - 2\gamma_0 \sum_{k=1}^{K} r_k - 2\gamma_1 \sum_{k=1}^{K} r_k x_{k|K} + 2\gamma_0\gamma_1 \sum_{k=1}^{K} x_{k|K} \Bigg\}. \quad (86)$$

The update for σ_w^2 may be obtained likewise. Again, in the case of an MPP observation, the sensor noise variance is calculated by performing the corresponding summations over the reduced set of indices at which $m_k = 1$.

M-step Update for σ_ε^2

Let

$$Q_4 = \frac{-K}{2}\log\left(2\pi\sigma_\varepsilon^2\right) - \sum_{k=1}^{K} \frac{\mathbb{E}\Big[(x_k - \rho x_{k-1} - \alpha I_k)^2\Big]}{2\sigma_\varepsilon^2}$$
$$= \frac{-K}{2}\log\left(2\pi\sigma_\varepsilon^2\right) - \sum_{k=2}^{K} \frac{\mathbb{E}\Big[(x_k - \rho x_{k-1} - \alpha I_k)^2\Big]}{2\sigma_\varepsilon^2} - \frac{\mathbb{E}\Big[(\alpha I_1)^2\Big]}{2\sigma_\varepsilon^2} \quad (87)$$

denote the term in Q containing σ_ε^2. Setting the partial derivative of Q_4 with respect to σ_ε^2 to 0, we obtain

$$\frac{\partial Q_4}{\partial \sigma_\varepsilon^2} = \frac{-K}{2\sigma_\varepsilon^2} + \frac{1}{2\sigma_\varepsilon^4} \sum_{k=2}^{K} \mathbb{E}\Big[(x_k - \rho x_{k-1} - \alpha I_k)^2\Big] + \frac{(\alpha I)^2}{2\sigma_\varepsilon^4} = 0. \quad (88)$$

Therefore,

$$\sigma_\varepsilon^2 = \frac{1}{K} \sum_{k=2}^{K} \Big\{ \mathbb{E}\big[x_k^2\big] - 2\rho\mathbb{E}\big[x_k x_{k-1}\big] + \rho^2\mathbb{E}\big[x_{k-1}^2\big] - 2\alpha I_k \mathbb{E}\big[x_k\big] + 2\alpha\rho I_k \mathbb{E}\big[x_{k-1}\big] \Big\}$$
$$+ \frac{\alpha^2}{K} \sum_{k=1}^{K} I_k^2$$
$$= \frac{1}{K}\Bigg\{ \sum_{k=2}^{K} U_k - 2\rho \sum_{k=1}^{K-1} U_{k,k+1} + \rho^2 \sum_{k=1}^{K-1} U_k - 2\alpha \sum_{k=2}^{K} I_k x_{k|K} + 2\alpha\rho \sum_{k=2}^{K} I_k x_{k-1|K}$$
$$+ \alpha^2 \sum_{k=1}^{K} I_k^2 \Bigg\}. \quad (89)$$

M-step Updates for β_0 and β_1

Let

$$Q_5 = \sum_{k=1}^{K} \mathbb{E}\Big[m_k(\beta_0 + \beta_1 x_k) - \log\left(1 + e^{\beta_0+\beta_1 x_k}\right)\Big] \tag{90}$$

denote the expectation term containing β_0 and β_1. Performing a Taylor expansion of the logarithm term around $x_{k|K}$ [46], we obtain

$$\begin{aligned}\log\left(1 + e^{\beta_0+\beta_1 x_k}\right) \approx \log\left(1 + e^{\beta_0+\beta_1 x_{k|K}}\right) + \beta_1 p_{k|K}(x_k - x_{k|K}) \\ + \frac{\beta_1^2}{2} p_{k|K}(1 - p_{k|K})(x_k - x_{k|K})^2.\end{aligned} \tag{91}$$

Taking the expected value on both sides yields

$$\begin{aligned}\mathbb{E}\Big[\log\left(1 + e^{\beta_0+\beta_1 x_k}\right)\Big] \approx &\log\left(1 + e^{\beta_0+\beta_1 x_{k|K}}\right) + \beta_1 p_{k|K}\mathbb{E}\big[x_k - x_{k|K}\big] + \frac{\beta_1^2}{2} p_{k|K}(1 - p_{k|K}) \\ &\times \mathbb{E}\big[(x_k - x_{k|K})^2\big] \\ = &\log\left(1 + e^{\beta_0+\beta_1 x_{k|K}}\right) + 0 + \frac{\beta_1^2}{2} p_{k|K}(1 - p_{k|K})\sigma_{k|K}^2.\end{aligned} \tag{92}$$

Therefore,

$$Q_5 \approx \sum_{k=1}^{K} \left[m_k(\beta_0 + \beta_1 x_{k|K}) - \log\left(1 + e^{\beta_0+\beta_1 x_{k|K}}\right) - \frac{\beta_1^2}{2} p_{k|K}(1 - p_{k|K})\sigma_{k|K}^2 \right]. \tag{93}$$

Now,

$$\begin{aligned}\frac{\partial p_{k|K}}{\partial \beta_0} &= \frac{\partial}{\partial \beta_0}\left[\frac{1}{1 + e^{-(\beta_0+\beta_1 x_{k|K})}}\right] = \frac{(-1)}{\left[1 + e^{-(\beta_0+\beta_1 x_{k|K})}\right]^2} \times \Big[- e^{-(\beta_0+\beta_1 x_{k|K})}\Big] \\ &= p_{k|K}(1 - p_{k|K}).\end{aligned} \tag{94}$$

And similarly,

$$\frac{\partial p_{k|K}}{\partial \beta_1} = p_{k|K}(1 - p_{k|K})x_{k|K}. \tag{95}$$

The partial derivative of Q_5 with respect to β_0 is

$$\frac{\partial Q_5}{\partial \beta_0} = \sum_{k=1}^{K} \left\{ m_k - p_{k|K} - \frac{\beta_1^2 \sigma_{k|K}^2}{2} \frac{\partial}{\partial \beta_0} \Big[p_{k|K}(1 - p_{k|K}) \Big] \right\}. \tag{96}$$

Setting this to 0, we obtain

$$0 = \sum_{k=1}^{K} \left[m_k - p_{k|K} - \frac{\beta_1^2 \sigma_{k|K}^2}{2} (1 - p_{k|K})(1 - 2p_{k|K}) p_{k|K} \right]. \tag{97}$$

And similarly for β_1 we have,

$$\frac{\partial Q_5}{\partial \beta_1} = \sum_{k=1}^{K} \left[m_k x_{k|K} - x_{k|K} p_{k|K} - \frac{\beta_1 \sigma_{k|K}^2}{2} p_{k|K}(1 - p_{k|K}) \big[2 + \beta_1 x_{k|K}(1 - 2p_{k|K}) \big] \right] = 0. \tag{98}$$

These two simultaneous equations involving the partial derivatives of Q_5 with respect to β_0 and β_1 can be solved numerically to obtain the corresponding M-step updates.

Approximation for the Expectation Term Containing $\lambda_{k,j}$

Let

$$Q_6 = \sum_{k=1}^{K} \sum_{j=1}^{J} \mathbb{E}\Big[\log(\lambda_{k,j}\Delta) n_{k,j} - \lambda_{k,j}\Delta \Big] \tag{99}$$

denote the expectation term containing $\lambda_{k,j}$. A Taylor expansion of the summed term around $x_{k|K}$ [46] yields

$$\begin{aligned} \log(\lambda_{k,j}\Delta) n_{k,j} - \lambda_{k,j}\Delta \approx & \log(\lambda_{k,j|K}\Delta) n_{k,j} - \lambda_{k,j|K}\Delta \\ & + \frac{1}{\lambda_{k,j|K}} \frac{\partial \lambda_{k,j|K}}{\partial x_k} (n_{k,j} - \lambda_{k,j|K}\Delta)(x_k - x_{k|K}) \end{aligned}$$

$$+\frac{1}{2}\left[\frac{1}{\lambda_{k,j|K}}\frac{\partial^2\lambda_{k,j|K}}{\partial x_k^2}(n_{k,j}-\lambda_{k,j|K}\Delta)-\frac{n_{k,j}}{\lambda_{k,j|K}^2}\left(\frac{\partial\lambda_{k,j|K}}{\partial x_k}\right)^2\right]$$
$$\times(x_k-x_{k|K})^2. \tag{100}$$

Taking the expected value on both sides, we obtain

$$\begin{aligned}\mathbb{E}\big[\log(\lambda_{k,j}\Delta)n_{k,j}-\lambda_{k,j|K}\Delta\big] \approx& \log(\lambda_{k,j|K}\Delta)n_{k,j}-\lambda_{k,j|K}\Delta\\ &+\frac{1}{\lambda_{k,j|K}}\frac{\partial\lambda_{k,j|K}}{\partial x_k}(n_{k,j}-\lambda_{k,j|K}\Delta)\mathbb{E}\big[x_k-x_{k|K}\big]\\ &+\frac{1}{2}\left[\frac{1}{\lambda_{k,j|K}}\frac{\partial^2\lambda_{k,j|K}}{\partial x_k^2}(n_{k,j}-\lambda_{k,j|K}\Delta)-\frac{n_{k,j}}{\lambda_{k,j|K}^2}\left(\frac{\partial\lambda_{k,j|K}}{\partial x_k}\right)^2\right]\\ &\times\mathbb{E}\big[x_k-x_{k|K}\big]^2\\ \approx& \log(\lambda_{k,j|K}\Delta)n_{k,j}-\lambda_{k,j|K}\Delta+0\\ &+\frac{1}{2}\left[\frac{1}{\lambda_{k,j|K}}\frac{\partial^2\lambda_{k,j|K}}{\partial x_k^2}(n_{k,j}-\lambda_{k,j|K}\Delta)-\frac{n_{k,j}}{\lambda_{k,j|K}^2}\left(\frac{\partial\lambda_{k,j|K}}{\partial x_k}\right)^2\right]\\ &\times\sigma_{k|K}^2. \end{aligned}\tag{101}$$

Therefore,

$$\begin{aligned}Q_6 \approx& \sum_{k=1}^{K}\sum_{j=1}^{J}\log(\lambda_{k,j|K}\Delta)n_{k,j}-\lambda_{k,j|K}\Delta\\ &+\frac{1}{2}\left[\frac{1}{\lambda_{k,j|K}}\frac{\partial^2\lambda_{k,j|K}}{\partial x_k^2}(n_{k,j}-\lambda_{k,j|K}\Delta)-\frac{n_{k,j}}{\lambda_{k,j|K}^2}\left(\frac{\partial\lambda_{k,j|K}}{\partial x_k}\right)^2\right]\sigma_{k|K}^2. \end{aligned}\tag{102}$$

This term can be numerically optimized with respect to the coefficient(s) in $\lambda_{k,j}$ that need to be determined.

3.5 Results and Discussion

3.5.1 Simulated Data

Only the extended model containing r_k and s_k was evaluated on simulated data. The procedure for simulating data is similar to that described in the earlier chapter (we set $\beta_1 = 1$ and calculated β_0 empirically on simulated data). We generated two datasets using a true $p_0 = 0.01$, but where the approximation $\hat{p}_0$ was slightly lower than and slightly higher than 0.01 in the actual data. We also set the indicator function $I_k = 1$ at 25 arbitrary locations.

The model parameters used and their estimates and shown in Table 3. Fig. 14 shows the estimation results in one of the cases. In the figure, the sub-panels respectively depict: (a) the probability of binary event occurrence p_k (blue) and its estimate (red, the green and black dots on the upper part of the sub-panel denote the presence or absence of the binary events); (b) the first continuous-valued variable r_k (blue) and its estimate (red); (c) the second continuous-valued variable s_k (blue) and its estimate (red); (d) the state x_k (blue) and its estimate (red, the cyan and black dots on the lower part of the sub-panel denote the presence or absence of binary inputs); (e) the sequence of RR-intervals $\mathrm{rr_i}$ (orange dots) and the estimated RR-interval mean μ (blue); (f) QQ plot for the residual errors of x_k. In general, there are good fits to r_k, s_k and the RR-intervals. However, the fits to p_k and x_k are better in one case ($\hat{p}_0 > 0.01$) than in the other. It is likely that since β_0 and β_1 (calculated using $\hat{p}_0$) appear in *exponent* terms, state estimation is more sensitive to them.

3.5.2 Experimental Data – Fear Conditioning Dataset

We evaluated the performance of the extended model containing r_k and s_k on the Fear Conditioning Dataset (PsPM-TC Dataset). We set $I_k = 1$ corresponding to the times at which the CS+, CS- and US stimuli were presented. We investigated model orders up to $q = 8$ based on partial autocorrelation plots of the RR-intervals similar to [103] and tried out values for η belonging to the set $\{-10^{-6}, -10^{-5}, -10^{-4}, -10^{-3}, -10^{-2}, -10^{-1}\}$ in the

Table 3: Parameter estimation with simulated data – one binary, two continuous and one spiking-type observation

True value	Estimate ($\hat{p}_0 > 0.01$)	Estimate ($\hat{p}_0 < 0.01$)
$\alpha = 0.04$	0.0082	0.0281
$\rho = 0.995$	0.9941	0.9947
$\delta_0 = -0.7$	-0.739	-0.704
$\delta_1 = 0.2$	0.2532	0.4417
$\sigma_w^2 = 0.003$	0.003	0.0031
$\gamma_0 = 0.35$	0.2716	0.3422
$\gamma_1 = 0.4$	0.5057	0.8848
$\sigma_v^2 = 0.002$	0.0019	0.0022
$\beta_0 = -4.5951$	-4.5458	-4.7015
$\beta_1 = 1$	1 (set)	1 (set)
$\sigma_\varepsilon^2 = 0.03$	0.0188	0.0058
$\theta_0 = 0.27432$	0.24028	0.29978
$\theta_1 = 0.83697$	0.83787	0.76209
$\theta_2 = -0.10511$	-0.07417	-0.05446
$\theta_3 = 234.22144$	239.28030	237.89351
$\eta = -0.005$	-0.001	-0.001

EM algorithm. Overfitting constraints similar to those described in the previous chapter also had to be applied. We also calculated β_1 and β_2 at the M-step (instead of setting β_0 = 1 and calculating β_0 empirically) as overfitting, rather than convergence, is the major concern with experimental data. We also included an additional constraint to prevent the α coefficient from becoming negative during estimation as this would imply that the external stimuli (for instance, the electric shock) decreases sympathetic arousal.

Pavlovian fear conditioning experiments have often sought to examine average differences in physiological features between the trial types. Since this is similar to the study of event-related potentials (ERP) with EEG data, we provide the ERP-like images for the three types of trials—CS-, CS+ without the US (CS+US-) and CS+ with the US (CS+US+). The state estimation results are shown in Fig. 15 for two representative cases. In both sub-figures, the sub-panels respectively depict: (a) the skin conductance signal z_k; (b) the probability of impulse occurrence (the green and black dots on the upper part of the sub-panel depict the detected SCR peak locations); (c) the phasic-derived value (solid green

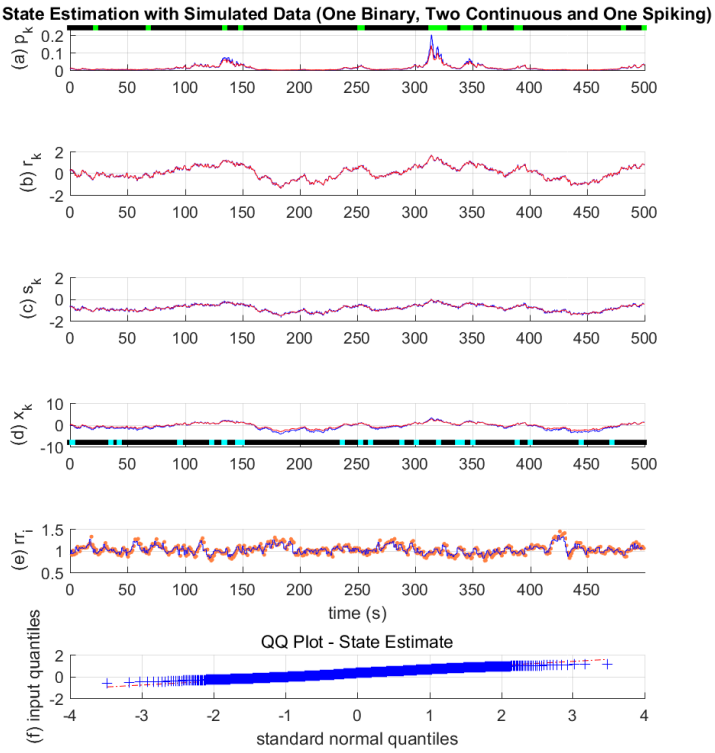

Figure 14: State estimation on simulated data for model with one binary, two continuous and one spiking-type observation.

line) and its fit (dotted line); (d) the tonic level (solid light mauve line) and its fit (dotted line); (e) the arousal state estimate (the cyan and black dots on the lower part of the sub-panel depict the stimuli input locations); (f) the sequence of RR-intervals rr_i (orange dots) and the estimated RR-interval mean μ (blue); (g) a 10 s ERP-like skin conductance plot for the CS- (green), CS+ without a shock (mauve – CS+US-) and CS+ with the shock (red – CS+US+) trials; (h) 10 s ERP-like arousal state plots along with their confidence intervals. The estimation results for all participants can be found in [94].

We divided the participants into three categories based on their physiological responses and state estimates. Overall, both the skin conductance and estimated arousal states are highest in the CS+US+ trails. For each of the participants in the first category, the average

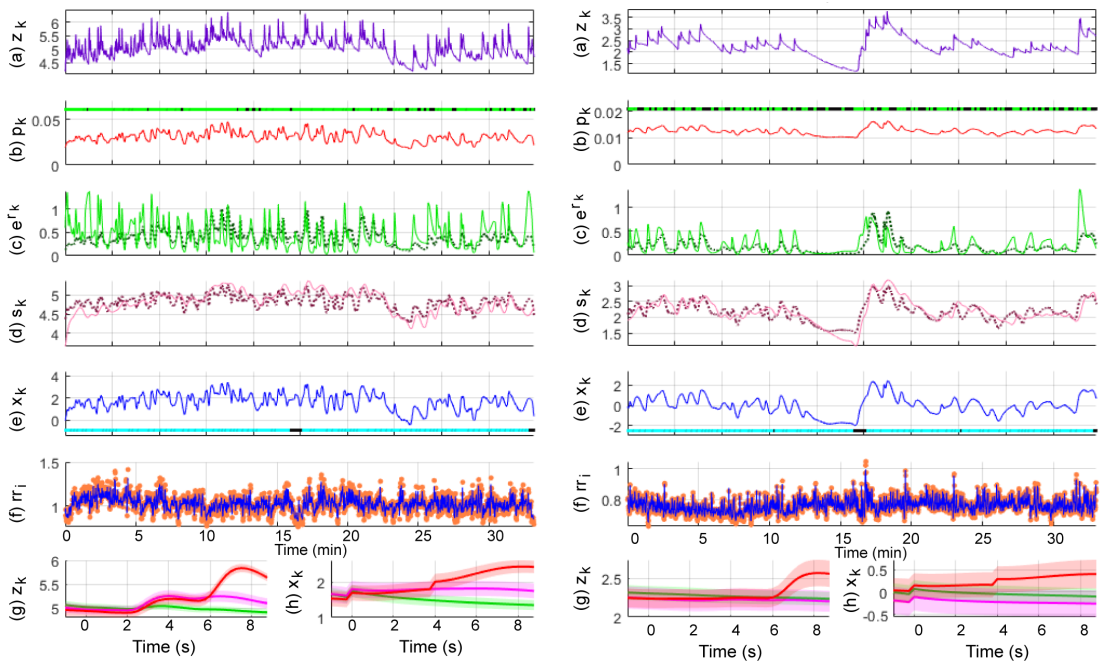

Figure 15: Sympathetic arousal estimation in the PsPM TC Dataset.

response to the CS+US+ is highest followed by the CS+US-. The CS- trials have the lowest average response. This is as expected. Clear gaps are visible between the averaged responses for each of the three types of trials. The gaps are visible for both averaged skin conductance and arousal. For participants in the next category, the gap between the CS+US- and CS- trials is very small. However, the average response for the CS+US+ is still larger. In the third category are participants for whom the general trend is that the averaged CS- curve tends to exceed the CS+US- curve. Participant 2 was an exception. Here, the averaged skin conductance and state estimates for the CS- and CS+US- were interchanged. Upon further analysis, it appeared that participant 2 developed a skin conductance arousal response to the CS- trials towards the end of the experiment. This was unusual as the participant should have by then learned that the CS- trials were never accompanied by the electric shock.

The KS plots (also provided in [94]) are close to the 45° diagonal for almost all the

participants indicating a good fit to the heartbeat observations. However, there are a few deviations from within the 95% confidence limits. The HDIG model developed by Barbieri et al., [98, 100] used time-varying θ_i coefficients that were estimated every 5 ms. The use of a fixed set of θ_i's estimated via maximum likelihood together with changes in arousal may have been insufficient to account for the HRV stochasticity completely.

In Pavlovian fear conditioning, a neutral stimulus is paired with an unpleasant stimulus such as an electric shock. Through repeated exposure, a subject learns an association between the two types of stimuli and eventually begins to display a response typically associated with the unpleasant stimuli to the neutral predictor as well. Due to ethical considerations involved in causing pain to human subjects, the intensity of the electric shocks used in fear conditioning experiments is often adjusted to be "highly annoying, but not painful" [108]. In the dataset used here, the shock intensity was adjusted to 90% of each subject's pain threshold [111]. Subjects have different pain thresholds and an electric shock that is not painful enough may not cause the subject to fear the US as much. This may be one of the possible reasons why variations are seen among the subjects analyzed in the PsPM TC Dataset. Ideally, we would expect to see the highest averaged responses (skin conductance and sympathetic arousal) for the CS+US+ trials and then for the CS+US- trials. The CS- trials would be expected to have the lowest averaged responses. However, this clear difference is only visible in participants in the first category. It appears that these participants learned the association between the CS+ cue and the electric shock and developed a fear response to the CS+ alone. In participants belonging to both the other categories, a clear separation with the average response for CS+US- being higher than the average response for CS- is not seen. There is almost no difference between the averaged responses for the CS+US- and CS- trials for participants in the second category and the response is inverted for participants in the third category. The reason for responses such as those seen for these participants is likely due to them not learning to fear the unpleasant electric shock enough. A further possibility for the lack of a response to the CS+ trials could

be the type of experiment that was used. The data come from a trace fear conditioning experiment. In trace fear conditioning, there is a gap between the time the CS+ ends and the application of the US. In delay fear conditioning, the CS+ stimuli co-terminate with the US without any time gap. Due to the closer pairing in time, the response to the CS+ stimuli in delay fear conditioning is usually larger than in trace fear conditioning [108]. Trace conditioning involves the hippocampus while delay conditioning predominantly involves the amygdala [108]. Finally, the experiment included the electric shock only in 50% of the CS+ trials. Therefore, a participant learns that not all CS+ trials precede a shock. The use of: (i) trace conditioning, (ii) CS+ only accompanying the US in 50% of the trials and (iii) shocks that may have not been unpleasant enough are possible reasons why only some participants had a response as expected.

It is also to be noted that there is a latency of a few seconds between the presentation of the stimuli and the appearance of an SCR (Fig. 15). This delay arises due to two primary reasons. One is related to the mechanism of sweat secretion [118] and the other is neural. Sweat largely comprises of water and sodium chloride (NaCl). When sympathetic nerve fibers stimulate the sweat glands, Cl^- ions enter the sweat ducts and cause Na^+ ions to enter as well due to the resulting electrochemical gradient. Potassium chloride (KCL) also enters the sweat ducts. The presence of these ions within the sweat ducts creates a second osmotic gradient causing water to enter through aquaporin channels. As a result, pressure builds up inside the sweat ducts and the liquid is finally forced out through the pores onto the skin surface. It is the presence of this salty liquid that causes the conductivity of the skin to drop. This mechanism of sweat secretion accounts for part of the delay from stimuli presentation to SCR appearance. There is, however, a second neural component to this delay. This neural aspect of SCR delay was the subject of investigation over several decades in the previous century. In particular, SCR latency was found to vary depending on phenomena such as the type of stimulus, habituation and the pairing of the stimuli. Some of these findings are described in [119, 120, 121]. Since the complete pathway from sensory

input to SCR appearance consists of the cerebral cortex, the hypothalamus, the spinal cord and sympathetic nerve fibers, the neural processing of the stimuli is thought to contribute to SCR latency as well.

3.5.3 Experimental Data – Smart Reasoning for Well-being at Home and at Work Dataset

We evaluated the performance of the reduced state-space model on the SWELL Dataset. Here we investigated model orders $q = 1, 2, 3, 4$ and selected $q = 3$ as the one with the smallest KS distance across all three trial blocks for the subject. We also tried out values for η ranging from $(-1) \times 10^{-5}$ to $(-1) \times 10^{-4}$ in increments of 10^{-5}. Estimates for two of the three experimental blocks are shown in Fig. 16. In both sub-figures, the sub-panels respectively depict: (i) the skin conductance signal z_k; (ii) the occurrence of the SCRs (binary); (iii) the arousal state estimate and its 95% confidence limits; (iv) the probability of SCR occurrence and its 95% confidence limits (the dotted line indicates the baseline probability); (v) the sequence of RR-intervals $\mathrm{rr_i}$ (orange dots) and the estimated RR-interval mean μ (blue); (vi) the certainty level (HAI). The estimates for the third experimental block and the KS plots can be found in [93]. The certainty level of arousal is calculated as the probability that p_k exceeds its baseline value. Sympathetic arousal as measured by the certainty level is highest at the start of the experiment in block UT and gradually decreases. Here, certainty exceeds 90% for about the first 20 minutes. In block RT, the certainty level does not exceed 90% even once, although it is highest at the start of the block. Arousal remains lowest during block CT and a gradual increase can be seen as the experiment progresses. This may be expected as the subject had to complete the writing tasks within the shortened time limit. The raw skin conductance signal can also be seen to decrease from block UT, to RT, to CT indicating a decline in arousal as the experiment proceeds.

The KS plots do not fall entirely within the 95% confidence limits although there is close agreement. In general, the better the CIF for modeling heart rate, the closer will be

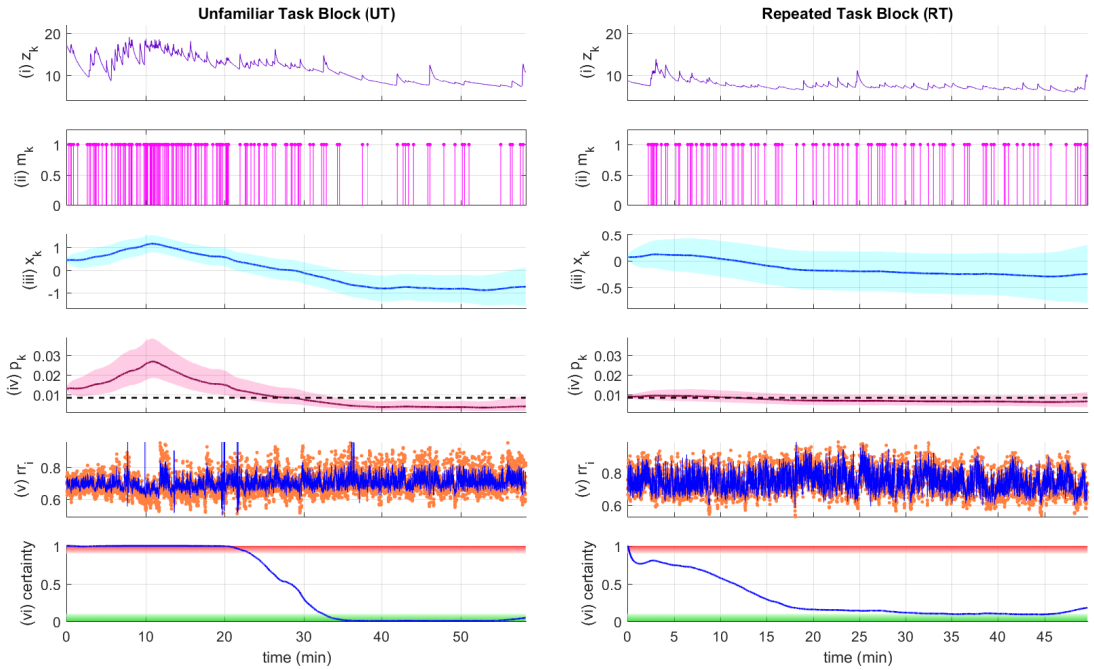

Figure 16: Sympathetic arousal estimation in the SWELL Dataset.

the plot to the 45° diagonal. The subject's EKG appeared to show some irregularity and this may have caused deviations from the HDIG model fits.

The SWELL experiment in [105] was originally designed to examine psychological stress in the presence of external interruptions and time restriction pressures. While interruptions and time constraints do cause anxiety, our analysis seems to point out that *unfamiliarity* or how new the task is seems to have caused the most stress for this particular subject. It may be the case that unfamiliarity with a task dominates the anxiety caused by both time constraints and unnecessary interruptions.

We also investigated whether the subject's typing correlated with the sympathetic arousal estimates. Unfortunately, this was complicated by certain factors. For instance, the key typing data did not agree with the final reports that the subject created. Since the subject was permitted to browse the web during the experiment, it may have been the case that the subject copied and pasted some text from the internet onto the reports. Thus, an

analysis of the key presses alone would be unable to account for a "working state" during the tasks. Moreover, based on the extracted timestamps, it also appeared that the subject appeared to wait for a while prior to beginning typing. Sympathetic arousal estimates tended to be high during the beginning of the sessions and it would have been interesting to investigate whether the subject made more mistakes (measured by the number of times the delete or backspace keys were pressed) during these times. Unfortunately, this was not possible due to the time interval between task commencement and when the subject began typing. Further analysis of the typing data, however, is a possible direction of future work. For instance, the typed key presses and the *backspace/delete key presses* may be considered as a sequence of binary correct/incorrect responses and a working state could thereby be estimated for further mental state analysis.

3.5.4 Model Parameter Selection

Recall that we select the θ_i coefficients and η separately. While this procedure eases computation (the θ_i's no longer have to be repeatedly estimated at the M-step until convergence), it also creates the challenge of having to optimize both types of parameters simultaneously. To illustrate, the θ_i coefficients are calculated offline via maximum likelihood. This may give rise to a KS plot indicating a reasonably good fit to the heartbeat observations. However, the inclusion of ηx_k into the HDIG mean during state estimation alters the KS plot and the KS distance. Consequently η may have to be finally selected based on a trade-off of maximizing $\bar{Q}$ subject to the KS plot remaining within or close to the 95% confidence bounds. Moreover, the separate selection of the θ_i's and η can also give rise to numerical issues; this can especially occur at larger η values. As the HDIG CIF involves derivatives and integrals over small numbers, the Newton-Raphson method used to solve for the state update can go into infeasible regions. The HDIG model we use here is computationally demanding. The use of a simpler probability density function to model RR-intervals may permit parameters related to heart rate to be estimated simultaneously

at the M-step. This would partly eliminate issues arising due to the need to separately optimize model parameters.

Alternatively, the θ_i coefficients could be considered as additional states. This would permit the simultaneous estimation of sympathetic arousal and the θ_i's. Additionally, it would also allow the θ_i's to be time-varying and account for some of the HRV stochasticity. In this case, the state vector to be estimated would consist of both x_k and the time-varying θ_i's. Barbieri et al., estimated a vector consisting of only the θ_i coefficients via Bayesian filtering in [100]. Their model could be extended to include x_k. However, the estimation of a vector state would also add to the complexity of the model and therefore other trade-offs may have to be considered as well.

4 Energy Production Estimation using Serum Cortisol Data

4.1 Hypothalamic-pituitary-adrenal Axis Function, Cortisol and Energy Production

Hormones play a critical role in regulating the body's internal environment. Cortisol, a hormone categorized as a glucocorticoid, is the body's primary stress hormone. Cortisol is secreted as the end result of a cascade of other secretory events in the HPA axis [92]. When the brain interprets sensory inputs as requiring cortisol secretion, the hypothalamus first begins to secrete corticotropin releasing hormone (CRH). This in turn causes the secretion of adrenocorticotropic hormone (ACTH) from the anterior pituitary. The secretion of ACTH triggers cortisol secretion from the adrenal glands. The secretion of cortisol has a negative feedback effect preventing the further secretion of CRH and ACTH. One of the major functions of cortisol is to raise blood glucose levels and give the body more energy in response to external stressors [122]. Most cortisol disorders involve either too much (*hyper*cortisolism) or too little (*hypo*cortisolism) of it in the bloodstream. Cortisol is secreted in pulses, and between 15–22 pulsatile secretions typically occur each day in a healthy adult [28]. Since cortisol is directly related to energy production, diminished levels of cortisol can can cause fatigue (e.g., in Addison's disease). In this chapter, we describe two different state-space models for estimating energy production based on blood cortisol measurements. The estimators may be eventually embedded within automated infusion pumps for treating cortisol disorders. However, a cortisol sensor that would be crucial for the functioning of such infusion pumps is presently lacking in the commercial market. Preliminary results of cortisol sensors described in the literature nevertheless provide the hope that such sensors will indeed be available in the near future [123, 124].

State-space models governing cortisol secretion dynamics have been developed previously (e.g., [125]). However, none of them explicitly take into account the pulsatile secretory nature *and* relate cortisol to the underlying state variable that the human body is actually

attempting to maintain within a desirable range. Here we describe two state-space models relating cortisol secretions to the body's internal energy state. The energy state is related to the pulsatile secretory events and to other continuous-valued phenomena related to blood cortisol. As before, we use Bayesian filtering within an EM framework for state estimation and model parameter recovery.

4.2 Data

4.2.1 Cushing's Dataset

In the absence of experimental data spanning multiple days, we first simulated cortisol measurements using the statistical models described in [30, 126] for a healthy subject and two Cushing's disease patients. Cushing's disease is a type of hypercortisolism that can be triggered by tumors or prolonged drug use [127].

Following [30], for a healthy subject, we draw pulse inter-arrival times for cortisol from a Gamma distribution with parameters $\alpha = 54$ and $\beta = 39$. These parameters are for inter-arrivals in terms of hours and we converted them to minutes during simulation. The pulse amplitudes h_k follow a time-of-day-dependent Gaussian distribution $h_k \sim \mathcal{N}(\mu_k, \kappa_k^2)$, where

$$\mu_k = 6.1 + 3.93 \sin\left(\frac{2\pi k}{1440}\right) - 4.75 \cos\left(\frac{2\pi k}{1440}\right) - 2.53 \sin\left(\frac{4\pi k}{1440}\right) - 3.76 \cos\left(\frac{4\pi k}{1440}\right), \quad (103)$$

$\kappa_k = \lambda\sqrt{\mu_k}$ and $\lambda = 0.1$ [30]. Given a vector of pulse input timings and amplitudes, we follow the solution of the coupled differential equations regulating the secretion of cortisol described in [128, 129] to obtain the serum cortisol profile over five days. We also use cortisol infusion and clearance rates of 0.0751 min^{-1} and 0.0086 min^{-1} based on the median rate parameters in [28] extracted for ten healthy subjects. Noise with standard deviation 0.5 μgdL^{-1} was finally added to the simulated observations.

Lee et al., [126] suggest inter-arrival times of 59 $\pm$ 11 min with amplitudes of 38 $\pm$ 2.5

$\mu gdL^{-1}min^{-1}$ for simulating the serum cortisol profile of a male Cushing's disease patient recruited in the study originally described in [130]. We used the Gamma and Gaussian distribution parameters corresponding to the Cushing's mean and standard deviation values above to simulate a second set of inter-arrival times and amplitudes respectively for a patient.

Berg et al., [130] noted the preservation of a cortisol circadian rhythm in a few Cushing's disease patients. We therefore simulated data for another hypothetical patient, one whose circadian rhythm was preserved by means of the time-of-day-dependent Gaussian pulse amplitudes. We set $\lambda = 2.5/\sqrt{38}$ when calculating the new pulse-amplitude standard deviation. We selected

$$\mu_k = 38.5 + 1.93 \sin\left(\frac{2\pi k}{1440}\right) - 1.6 \cos\left(\frac{2\pi k}{1440}\right) - 1.5 \sin\left(\frac{4\pi k}{1440}\right) - 3.5 \cos\left(\frac{4\pi k}{1440}\right) \quad (104)$$

to produce amplitudes in the same approximate range as for the first Cushing's disease patient with the same Gamma inter-arrival distribution and simulated a third set of measurements.

The cortisol measurements from these three subjects form the Cushing's Dataset.

4.2.2 Chronic Fatigue Syndrome and Fibromyalgia Syndrome Dataset

We secondly used an experimental dataset. The blood cortisol measurements in this dataset are from the study described in [131] conducted at the University of Michigan Medical Center. The study sought to investigate HPA axis function in patients diagnosed with fibromyalgia syndrome (FMS), chronic fatigue syndrome (CFS) or both. Blood samples were drawn intravenously at 10 min intervals over a 24 h period from patients and matched controls and were later processed to yield cortisol concentrations. Pednekar et al., [132] deconvolved the blood cortisol measurements from 31 of the 36 data record pairs that were available. Deconvolution yields the cortisol pulse amplitudes and timings as well as

the cortisol infusion and clearance rates with which a minute-by-minute cortisol profile can be reconstructed. We utilize the cortisol pulses and the reconstructed blood cortisol measurements for energy state estimation. Thirteen of the 31 patients were diagnosed with CFS (we label this the CFS patient group) and the remaining 18 were diagnosed with FMS or both FMS and CFS (we label this the FMS patient group). The data from these 31 subject pairs form the CFS and FMS Dataset.

4.3 Methods

4.3.1 State-space Model with One Binary and Two Continuous Observations

We first use a state-space model with one binary and two continuous observations for cortisol-related energy state estimation. The model is described in [133]. Unlike the earlier model in chapter 2, however, here we change the state equation to explicitly account for the circadian rhythmicity of cortisol [30]. We assume the energy state x_k evolves with time as

$$x_k = \rho x_{k-1} + I_k + \varepsilon_k, \tag{105}$$

where

$$I_k = \sum_{i=1}^{2} a_i \sin\Big(\frac{2\pi ik}{1440}\Big) + b_i \cos\Big(\frac{2\pi ik}{1440}\Big), \tag{106}$$

$\varepsilon_k \sim \mathcal{N}(0, \sigma_\varepsilon^2)$ is process noise and ρ is a coefficient to be determined along with the a_i and b_i terms. I_k acts as a forcing function in keeping with known energy variations during wakefulness and sleep in a 24 h period. It is generally taken that the circadian secretory pattern of ACTH imposes the same rhythm on cortisol [125].

We analyze our data at a time resolution of one minute (24 h = 1440 min). The presence or absence of a cortisol pulse each minute forms a binary point process and we assign $m_k = \{1, 0\}$ accordingly. The occurrence of a cortisol pulse at each time instant is a

Bernoulli distributed random variable with probability p_k. As in earlier cases, we take p_k to be related to x_k through

$$\log\left(\frac{p_k}{1-p_k}\right) = \beta_0 + \beta_1 x_k, \tag{107}$$

where β_0 and β_1 are coefficients to be determined. We also calculate the upper and lower envelopes of the reconstructed blood cortisol levels using MATLAB's *envelope* function to be used as the continuous-valued observations. Fig. 17 shows the simulated data for two of the subjects in the Cushing's Dataset. In both sub-figures, the sub-panels respectively depict: (i) the simulated blood cortisol concentrations; (ii) the pusaltile secretions; (iii) the upper (green) and lower (mauve) envelopes of the cortisol concentrations. Here, we use the *peak* option in MATLAB to estimate the envelopes as spline interpolations over local extrema. Labeling the upper and lower envelopes as r_k and s_k respectively, we assume the following linear relationships with x_k:

$$r_k = \gamma_0 + \gamma_1 x_k + v_k \tag{108}$$

and

$$s_k = \delta_0 + \delta_1 x_k + w_k. \tag{109}$$

Here, $v_k \sim \mathcal{N}(0, \sigma_v^2)$ and $w_k \sim \mathcal{N}(0, \sigma_w^2)$ represent sensor noise terms. The terms $\gamma_0, \gamma_1, \delta_0$ and δ_1 are regression coefficients to be determined.

A modification to the upper and lower envelopes r_k and s_k was necessary in the case of Cushing's disease where the levels of serum cortisol are much higher. This elevation is likely due to a malfunctioning of the feedback control mechanism governing the secretion of cortisol [127]. For a healthy subject, the serum cortisol levels can decrease to a value close to zero, although not to zero itself. Therefore, we can assume that zero forms a *lower*

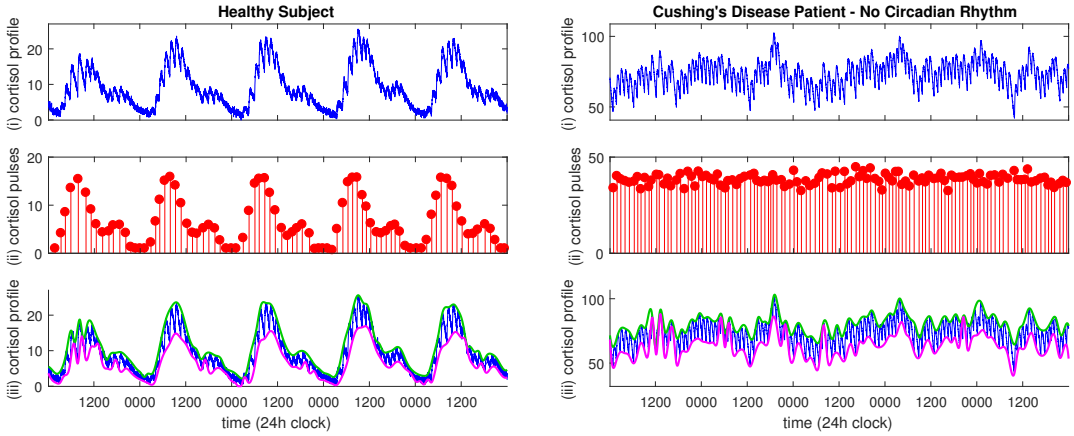

Figure 17: Simulated serum cortisol data for a healthy subject and a Cushing's patient.

baseline. We hypothesize that the control system malfunction in Cushing's is caused by the body establishing a *new* lower baseline that cortisol levels are not allowed to drop below. This lower baseline is much higher than normal and is likely caused by the body gradually developing a resistance to the chronically high levels of cortisol. For both the (simulated) Cushing's patients we take the minimum cortisol level across the five days as the lower baseline and subtract it from the cortisol profile to obtain the modified upper and lower envelopes.

The E-step is very similar to what was described for the state-space model with one binary and two continuous observations in chapter 2. The forward filter equations are as follows:

$$x_{k|k-1} = \rho x_{k-1|k-1} + I_k, \tag{110}$$

$$\sigma^2_{k|k-1} = \rho^2 \sigma^2_{k-1|k-1} + \sigma^2_\varepsilon, \tag{111}$$

$$C_k = \frac{\sigma^2_{k|k-1}}{\sigma^2_v \sigma^2_w + \sigma^2_{k|k-1}(\gamma_1^2 \sigma^2_w + \delta_1^2 \sigma^2_v)}, \tag{112}$$

$$\begin{aligned} x_{k|k} = x_{k|k-1} + C_k \Big[& \beta_1 \sigma^2_v \sigma^2_w (m_k - p_{k|k}) + \gamma_1 \sigma^2_w (r_k - \gamma_0 - \gamma_1 x_{k|k-1}) \\ & + \delta_1 \sigma^2_v (s_k - \delta_0 - \delta_1 x_{k|k-1}) \Big], \end{aligned} \tag{113}$$

and

$$\sigma^2_{k|k} = \left[\frac{1}{\sigma^2_{k|k-1}} + \beta_1^2 p_{k|k}(1 - p_{k|k}) + \frac{\gamma_1^2}{\sigma_v^2} + \frac{\delta_1^2}{\sigma_w^2} \right]^{-1}. \tag{114}$$

Again, the appearance of $p_{k|k}$ on both sides of the state update equation requires the use of Newton's method for solving. The model parameters are estimated at the M-step and the equations are similar to those for the other state-space methods. The a_i and b_i terms are estimated at the M-step by way of optimization [133].

4.3.2 State-space Model with One Marked Point Process and One Continuous Observation

The second state-space model considers the pulsatile profile to form a set of MPP observations (m_k, r_k) and the reconstructed minute-by-minute blood cortisol concentrations to form a continuous-valued variable s_k. Thus, the model is an extension of the earlier MPP-based approach described in chapter 2. In this model, we assume that: (i) x_k follows a random walk; (ii) the (m_k, r_k) pairs are related to x_k similar to eq. (30); (iii) s_k is linearly related to x_k similar to eq. (109). With this formulation, the filter equations switch between those in [44] and [72] depending on m_k (i.e., they switch between the model with one binary and one continuous observation and the model with one binary and two continuous observations).

The filter equations for this state-space model are as follows:

$$x_{k|k-1} = x_{k-1|k-1} \tag{115}$$

and

$$\sigma^2_{k|k-1} = \sigma^2_{k-1|k-1} + \sigma^2_{\varepsilon}. \tag{116}$$

If $m_k = 0$, we have

$$C_k = \frac{\sigma^2_{k|k-1}}{\gamma_1^2 \sigma^2_{k|k-1} + \sigma_v^2}, \tag{117}$$

$$x_{k|k} = x_{k|k-1} + C_k[\beta_1 \sigma_v^2 (m_k - p_{k|k}) + \gamma_1 (r_k - \gamma_0 - \gamma_1 x_{k|k-1})], \tag{118}$$

$$\sigma^2_{k|k} = \left[\frac{1}{\sigma^2_{k|k-1}} + \beta_1^2 p_{k|k}(1 - p_{k|k}) + \frac{\gamma_1^2}{\sigma_v^2} \right]^{-1}, \tag{119}$$

and if $m_k = 1$, we instead have

$$C_k = \frac{\sigma^2_{k|k-1}}{\sigma_v^2 \sigma_w^2 + \sigma^2_{k|k-1}(\gamma_1^2 \sigma_w^2 + \delta_1^2 \sigma_v^2)}, \tag{120}$$

$$\begin{aligned} x_{k|k} = x_{k|k-1} + C_k \Big[& \beta_1 \sigma_v^2 \sigma_w^2 (m_k - p_{k|k}) + \gamma_1 \sigma_w^2 (r_k - \gamma_0 - \gamma_1 x_{k|k-1}) \\ & + \delta_1 \sigma_v^2 (s_k - \delta_0 - \delta_1 x_{k|k-1}) \Big], \end{aligned} \tag{121}$$

and

$$\sigma^2_{k|k} = \left[\frac{1}{\sigma^2_{k|k-1}} + \beta_1^2 p_{k|k}(1 - p_{k|k}) + \frac{\gamma_1^2}{\sigma_v^2} + \frac{\delta_1^2}{\sigma_w^2} \right]^{-1}. \tag{122}$$

The M-step updates for the model parameters are similar to those for the other state-space models in earlier chapters. However, one notable change is made with experimental data. Now x_k can overfit to the continuous-valued blood cortisol concentration s_k. This overfitting can also occur very rapidly within the first few EM iterations. Moreover, the overfitting to a continuous variable is a problem that is *common to other point process Bayesian filters as well*. Therefore, we apply a two-part modification at the M-step to prevent this from occurring. First, we take $\theta = \{\delta_0, \delta_1, \sigma_w^2\}$ to be the model parameters governing x_k's fit to s_k, and modify the $(l+1)^{\text{th}}$ M-step update for θ to

$$\theta^{(l+1)} = \theta^{(l)} + \psi \Big[\theta^{(l+1)}_{\text{pred}} - \theta^{(l)} \Big], \tag{123}$$

where $0 < \psi \leq 1$. Setting $\psi = 1$ would mean that we perform the full update $\theta^{(l+1)} = \theta^{(l+1)}_{\text{pred}}$. Instead we choose $\psi < 1$ taking a smaller step in the direction of the next value. Secondly, the noise variance estimate σ_w^2 critically influences overfitting. Without any overfitting control, σ_w^2 drops rapidly within the first few EM iterations causing x_k to overfit to s_k. Therefore, as a second measure, we initialize $\sigma^2_{w,\text{init}}$ at a larger value and allow θ to only update until σ_w^2 reaches a threshold $\sigma^2_{w,\text{stop}}$. The other model parameters continue to update until convergence. This is a form of the early-stopping used in [72] but with the ψ term now included. Tuning $\sigma^2_{w,\text{init}}$, ψ and $\sigma^2_{w,\text{stop}}$ provides a fine-grained means of controlling the degree of overfitting to a desired level.

4.3.3 Derivation of E-step Filter Equations for State-space Model with One Marked Point Process and One Continuous Observation

At this point, we briefly digress again to show the E-step filter derivations of the model with one MPP and one continuous observation. Mathematically, this is the most complex form we will encounter with an MPP present (the E-step equations for the earlier MPP-based model in chapter 2 can be derived by dropping the continuous-valued variable here).

First consider the case when $m_k = 0$ and $p(m_k \cap r_k | x_k) = (1 - p_k) = e^{\log(1-p_k)}$. The posterior density is

$$\begin{aligned} p(x_k|y^k) &= \frac{p(x_k|y^{k-1})p(y_k|x_k)}{p(y_k|y^{k-1})} \\ &= \frac{p(x_k|y^{k-1})p(m_k \cap r_k|x_k)p(s_k|x_k)}{p(y_k|y^{k-1})} \\ &\propto \exp\left[\frac{-(x_k - x_{k|k-1})^2}{2\sigma^2_{k|k-1}} + \log(1-p_k) - \frac{(s_k - \delta_0 - \delta_1 x_k)^2}{2\sigma_w^2}\right]. \end{aligned} \tag{124}$$

Therefore,

$$q_1 = \log\Big[p(x_k|y^k)\Big] = \frac{-(x_k - x_{k|k-1})^2}{2\sigma^2_{k|k-1}} + \log(1-p_k) - \frac{(s_k - \delta_0 - \delta_1 x_k)^2}{2\sigma_w^2} + \text{const.} \tag{125}$$

We first take the partial derivative of q_1 and set it to 0. This yields

$$\frac{\partial q_1}{\partial x_k} = \frac{-2(x_k - x_{k|k-1})}{2\sigma^2_{k|k-1}} + \frac{1}{(1-p_k)}[-\beta_1 p_k(1-p_k)] + \frac{2\delta_1(s_k - \delta_0 - \delta_1 x_k)}{2\sigma^2_w} = 0. \tag{126}$$

Therefore,

$$-\beta_1 p_k + \frac{\delta_1(s_k - \delta_0 - \delta_1 x_k)}{\sigma^2_w} = \frac{(x_k - x_{k|k-1})}{\sigma^2_{k|k-1}}. \tag{127}$$

Since $m_k = 0$, we can write $-\beta_1 p_k = \beta_1(0 - p_k) = \beta_1(m_k - p_k)$. Therefore,

$$\beta_1(m_k - p_k) + \frac{\delta_1(s_k - \delta_0 - \delta_1 x_k)}{\sigma^2_w} = \frac{(x_k - x_{k|k-1})}{\sigma^2_{k|k-1}}. \tag{128}$$

Solving for x_k yields,

$$x_k = x_{k|k-1} + \frac{\sigma^2_{k|k-1}}{\delta_1^2\sigma^2_{k|k-1} + \sigma^2_w}\Big[\sigma^2_w\beta_1(m_k - p_k) + \delta_1(s_k - \delta_0 - \delta_1 x_{k|k-1})\Big], \tag{129}$$

which is the filter equation for the state x_k in [44]. Taking the second derivative yields

$$\frac{\partial^2 q_1}{\partial x_k^2} = \frac{-1}{\sigma^2_{k|k-1}} - \beta_1^2 p_k(1-p_k) - \frac{\delta_1^2}{\sigma^2_w}. \tag{130}$$

This provides the variance update in [44]

$$-\left[\frac{\partial^2 q_1}{\partial x_k^2}\right]^{-1} = \left[\frac{1}{\sigma^2_{k|k-1}} + \beta_1^2 p_{k|k}(1-p_{k|k}) + \frac{\delta_1^2}{\sigma^2_w}\right]^{-1}. \tag{131}$$

Next consider the case when $m_k = 1$ and

$$p(m_k \cap r_k|x_k) = e^{\log(p_k)}\frac{1}{\sqrt{2\pi\sigma_v^2}}e^{\frac{-(r_k - \gamma_0 - \gamma_1 x_k)^2}{2\sigma_v^2}}. \tag{132}$$

This yields the posterior

$$p(x_k|y_{1:k}) \propto \exp\left[\frac{-(x_k - x_{k|k-1})^2}{2\sigma_{k|k-1}^2} + \log(p_k) - \frac{(r_k - \gamma_0 - \gamma_1 x_k)^2}{2\sigma_v^2} - \frac{(s_k - \delta_0 - \delta_1 x_k)^2}{2\sigma_w^2}\right]. \tag{133}$$

Using the label q_2 in this case and taking the first partial derivative yields

$$\begin{aligned}\frac{\partial q_2}{\partial x_k} &= \frac{-(x_k - x_{k|k-1})}{\sigma_{k|k-1}^2} + \frac{1}{p_k}\beta_1 p_k(1-p_k) + \frac{\gamma_1(r_k - \gamma_0 - \gamma_1 x_k)}{\sigma_v^2} + \frac{\delta_1(s_k - \delta_0 - \delta_1 x_k)}{\sigma_w^2} \\ &= \frac{-(x_k - x_{k|k-1})}{\sigma_{k|k-1}^2} + \beta_1(1-p_k) + \frac{\gamma_1(r_k - \gamma_0 - \gamma_1 x_k)}{\sigma_v^2} + \frac{\delta_1(s_k - \delta_0 - \delta_1 x_k)}{\sigma_w^2}. \end{aligned} \tag{134}$$

Since $m_k = 1$, we may replace $(1 - p_k)$ with $(m_k - p_k)$. Setting the partial derivative to 0 and solving for x_k yields

$$\begin{aligned} x_k = \; & x_{k|k-1} + \frac{\sigma_{k|k-1}^2}{\sigma_v^2\sigma_w^2 + \sigma_{k|k-1}^2(\gamma_1^2\sigma_w^2 + \delta_1^2\sigma_v^2)} \\ & \times \Big[\sigma_v^2\sigma_w^2\beta_1(m_k - p_k) + \gamma_1\sigma_w^2(r_k - \gamma_0 - \gamma_1 x_{k|k-1}) + \delta_1\sigma_v^2(s_k - \delta_0 - \delta_1 x_{k|k-1})\Big]. \end{aligned} \tag{135}$$

Similarly, the second derivative is

$$\frac{\partial^2 q_2}{\partial x_k^2} = \frac{-1}{\sigma_{k|k-1}^2} - \beta_1^2 p_k(1-p_k) - \frac{\gamma_1^2}{\sigma_v^2} - \frac{\delta_1^2}{\sigma_w^2}. \tag{136}$$

Therefore, the variance update is

$$-\left[\frac{\partial^2 q_2}{\partial x_k^2}\right]^{-1} = \left[\frac{1}{\sigma_{k|k-1}^2} + \beta_1^2 p_{k|k}(1-p_{k|k}) + \frac{\gamma_1^2}{\sigma_v^2} + \frac{\delta_1^2}{\sigma_w^2}\right]^{-1}. \tag{137}$$

These update equations for the mean and variance when $m_k = 1$ correspond to those given in [72].

4.4 Results and Discussion

4.4.1 Simulated Data

We tested out the performance of the second state-space model on simulated data (the first state-space model with one binary and two continuous observations is almost identical to the model in chapter 2 that was earlier evaluated on simulated data). The procedure for generating synthetic data is very similar to the earlier cases. In this case, we simulate two sets of data with a true $p_0 = 0.05$, but where the approximation $\hat{p}_0$ is firstly slightly less than and secondly slightly greater than 0.05. Fig. 18 shows the state estimation results in one of the cases. In the figure, the sub-panels respectively depict: (a) the MPP observations (blue) and the estimated fit to r_k (red); (b) the probability of point process event occurrence p_k (blue) and its estimate (red); (c) the continuous variable s_k (blue) and its estimate (red); (d) the state x_k (blue) and its estimate (red); (e) the QQ plot for the residual errors of x_k. The model parameters used to generate both datasets and their final EM estimates are shown in Table 4.

Table 4: Parameter estimation with simulated data – one MPP and one continuous observation

True value	Estimate ($\hat{p}_0 < 0.05$)	Estimate ($\hat{p}_0 > 0.05$)
$\gamma_0 = 0.2$	0.18761	0.13222
$\gamma_1 = 0.7$	0.63388	0.77937
$\sigma_v^2 = 0.05$	0.04510	0.06064
$\delta_0 = -0.3$	-0.34946	-0.33654
$\delta_1 = 0.4$	0.36998	0.45076
$\sigma_w^2 = 0.002$	0.00195	0.00208
$\sigma_\varepsilon^2 = 0.005$	0.00644	0.00414

In general, there is a good agreement between the true values and their estimates. Since no constraint has been placed on the sign of the mark values, the algorithm can be used in other applications as well.

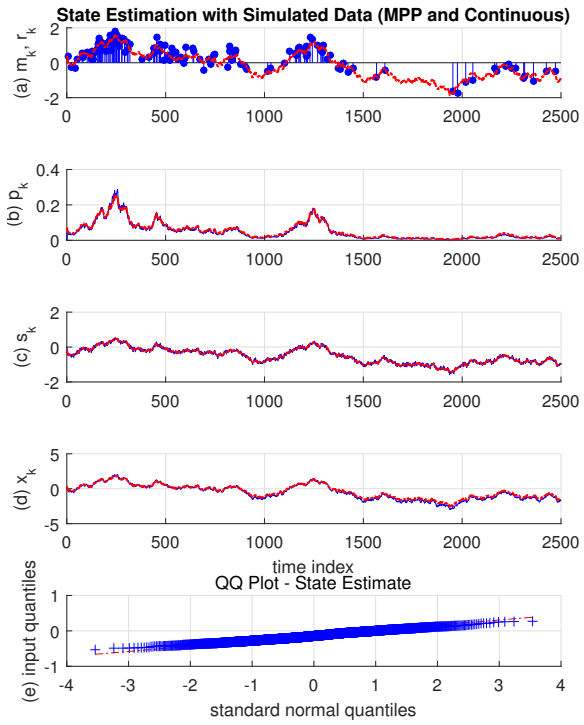

Figure 18: State estimation on simulated data for model with one MPP and one continuous observation.

4.4.2 Experimental Data – Cushing's Dataset (Simulated)

We tested out the first state-space model with one binary and two continuous observations on the Cushing's Dataset. Energy state estimates and fits to the observations for the healthy subject and the Cushing's patient without the circadian rhythm are shown in Fig. 19. In both sub-figures, the sub-panels respectively depict: (i) the raw blood cortisol data z_k; (ii) the pulsatile secretions m_k; (iii) the probability of pulse occurrence p_k; (iv) the upper envelope r_k (solid green) and its estimate (dashed); (v) the lower envelope s_k (solid mauve) and its estimate (dashed); (vi) the energy production state x_k. The results for the other Cushing's patient can be found in [133]. The healthy subject's energy state varies fairly consistently following circadian periodicity. A large energy peak typically appears between 6:00–10:00 a.m. as expected, and a slight drop occurs in the afternoon. A secondary peak

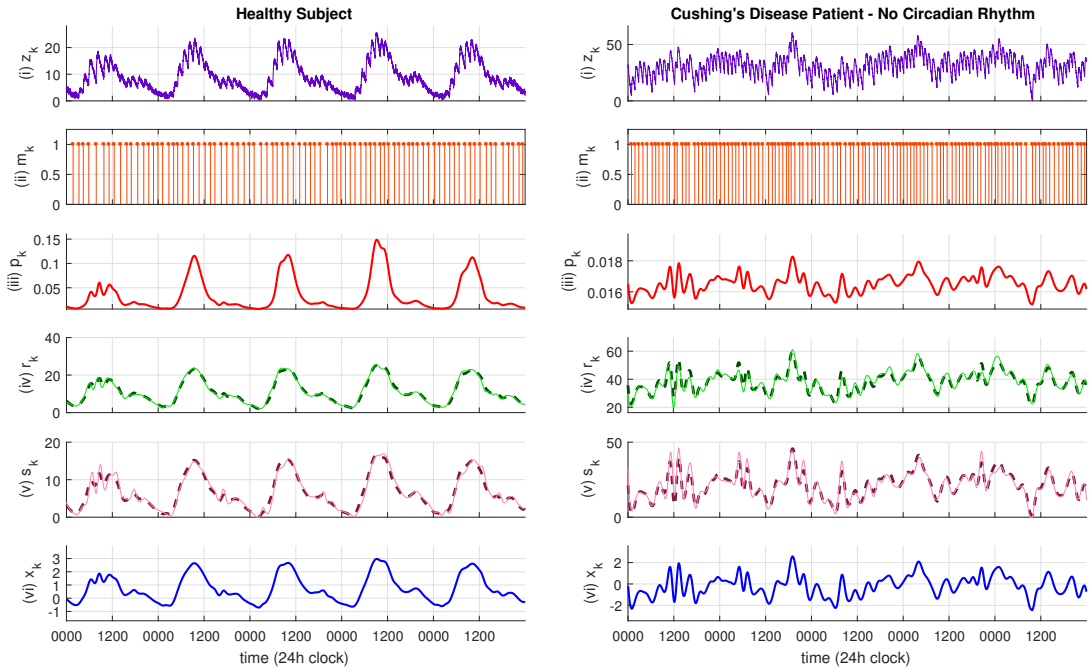

Figure 19: Cortisol-related energy state estimation on a simulated healthy subject and a Cushing's patient.

occurs between late afternoon and early evening.

No such circadian energy variation is seen for the first Cushing's disease patient. There are significant nighttime energy increases and daytime drops. This is expected as the simulations used a non-circadian probability distribution.

Interestingly, the second Cushing's disease patient also does not exhibit circadian energy variations despite the amplitudes being drawn from a Gaussian distribution with a circadian mean. We use coefficients for the sinusoids that do not give rise to a circadian rhythm with peaks as prominent as for the healthy subject. Moreover, the λ is also larger than 0.1. While the pulse amplitudes for this patient do follow a somewhat repetitive pattern, the serum cortisol levels do not exhibit that same pattern. It is likely that the Gamma-distributed inter-arrival times for Cushing's cause this circadian rhythm disruption.

A healthy subject regularly experiences more energy during hours of wakefulness than

during sleep. Recall that cortisol primarily raises blood glucose levels. One would expect therefore, that unusually high cortisol levels in a Cushing's disease patient constantly provide more energy. Instead, patients frequently experience fatigue and insomnia [134, 135]. While several factors could underlie this phenomena, we present a new perspective purely based on the energy state model. Our analysis offers the explanation that fatigue during the day and sleep disturbances at night could be due to the way in which energy varies during each 24 h period. For Cushing's patients, drops in energy are frequently seen in the daytime, as are increases during the night and likely cause daytime fatigue and nighttime sleeping difficulties.

4.4.3 Experimental Data – Chronic Fatigue Syndrome and Fibromyalgia Syndrome Dataset

We evaluated the second state-space model comprising of one MPP and one continuous observation on the CFS and FMS Dataset. Estimation results on two representative cases are shown in Fig. 20. Within each sub-figure, the sub-panels respectively depict: (a) the raw blood cortisol data z_k; (b) the deconvolved cortisol pulses (blue) and the estimated fit to r_k (red); (c) the reconstructed blood cortisol profile s_k (orange) and its estimated fit (red); (d) the probability of pulse occurrence p_k with its 95% confidence limits; (e) the energy production state x_k with its 95% confidence limits. We likewise estimated energy states for all patients and matched controls in the dataset. Typically, cortisol levels begin to rise early morning during late sleep and reach peak values within 30-45 min after awakening. Cortisol concentrations thereafter gradually decline and tend to reach nadir values at about midnight during sleep [136, 137]. In general, this pattern is visible for x_k in both patients and matched controls. However, we did not observe any major differences between *average* energy levels for the patients and the matched controls as shown in Fig. 21. In this figure, the sub-figures depict the mean pulse probabilities and estimated states for CFS and FMS patients and matched controls. Within each sub-figure, the sub-panels respectively depict:

(a) the mean probabilities for patients (red) and controls (blue); (b) the mean energy states for patients (red) and controls (blue). Care also must be taken when interpreting the average values since the raw x_k estimates can lie in slightly different ranges for different subjects.

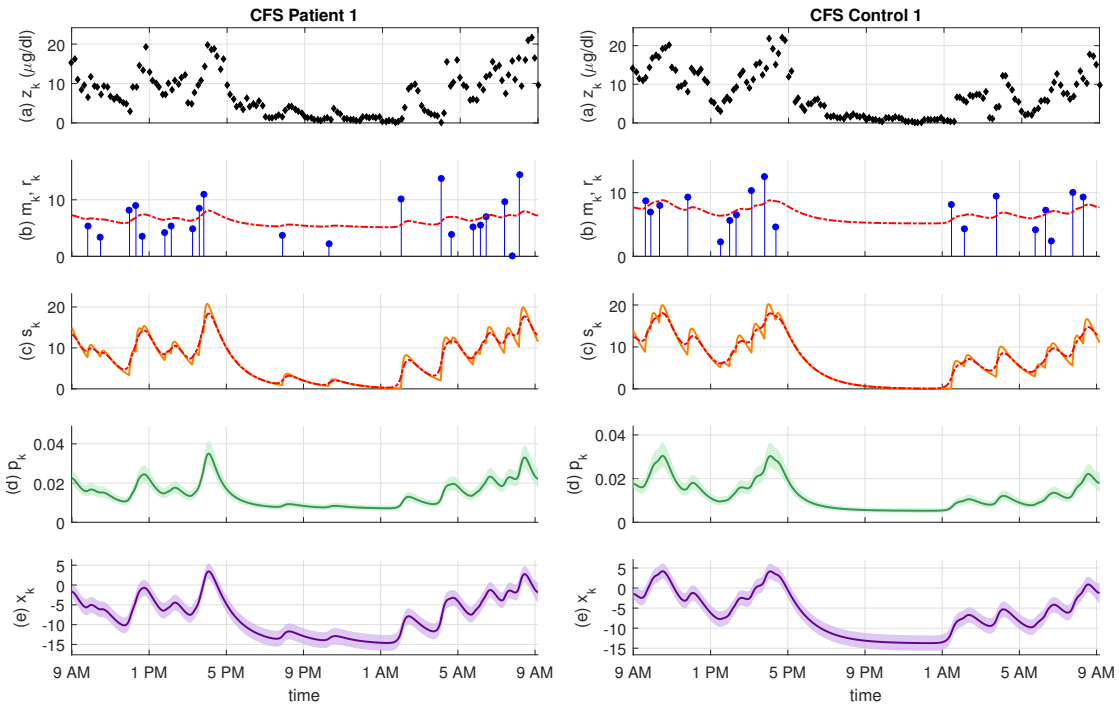

Figure 20: Cortisol-related energy state estimation on a CFS patient and the corresponding healthy control.

The CFS and FMS Data contains measurements from patients diagnosed with CFS, FMS or both, and matched controls. CFS is a chronic condition of unknown etiology characterized by debilitating fatigue [138]. It is also accompanied by symptoms such as difficulty concentrating, sleep disturbance, pain and post-extertional malaise [138]. FMS is a similar condition sharing certain symptoms with CFS.

There has been a lack of consensus in studies investigating cortisol variations in CFS patients and healthy controls. A review of the data prior to 2003 can be found in [139], and a more recent review of the studies after 2003 in [140]. Single sample studies, which have been criticized as being of limited use, have reported no change, increases and decreases

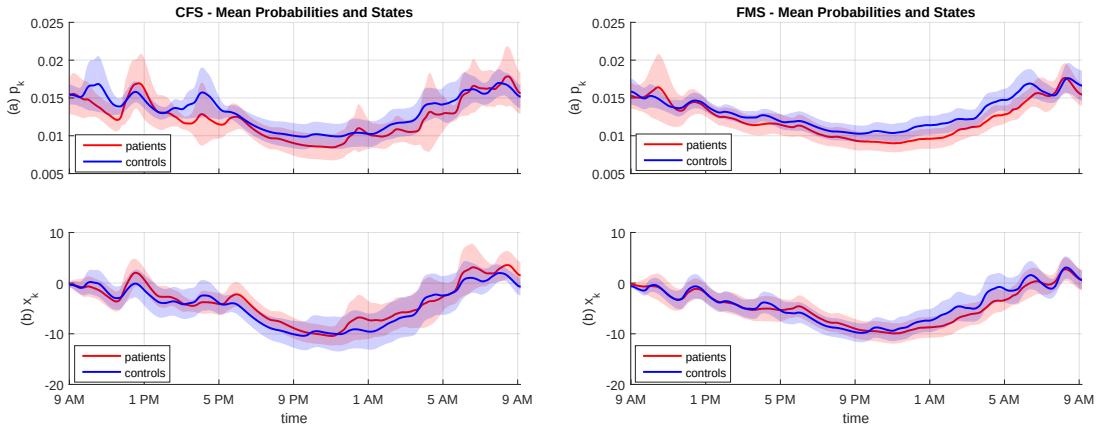

Figure 21: Mean pulse probabilities and energy states for CFS and FMS patients and controls.

between CFS patients and healthy controls [139, 140]. Two of the most important serial sample studies in [131] and [141] where blood samples where drawn every 10-15 min over a 24 h period found no significant differences between patients and controls at any point in time [140]. While evidence does point to hypocortisolism in at least some CFS patients, it is likely that there is no single cause for CFS [140]. Moreover, whether changes in HPA axis function are causal or consequential also remains unclear. Studies investigating cortisol changes between FMS patients and healthy subjects have also yielded mixed results [142, 143]. The serial sample study in [131] (we use this data here) also found no significant differences between FMS, and FMS and CFS patients and matched controls. This may be a possible reason as to why we also did not see any significant differences between the average values for patients and controls.

5 Machine Learning Approaches for State Estimation

5.1 Limitations with Traditional Approaches and the Need for Hybrid Estimation

We have thus far described several state-space models and corresponding EM algorithms for estimating unobserved quantities within the human body. There are, however, a few drawbacks with the general state-space EM approach. Firstly, the overall approach is somewhat inflexible and does not scale well with the number of observations (features). For instance, there is a significant increase in filter complexity as a state-space model goes from incorporating one binary observation to incorporating one binary *and* one continuous observation. Adding yet another continuous observation increases the complexity even further. Secondly, the models do not include a convenient means of incorporating external influences such as domain expertise or subject-provided labels. For instance, assume that a particular state x_k related to a hormone needs to decrease at meal times or at night. If this "knowledge" to be included in the model, the state equation has to be explicitly changed to incorporate it. Suppose now, that the external influences are also patient-specific. This would necessitate the development of an entire series of models, each with its own EM algorithm. The development of machine learning methods to perform state estimation [144, 145] now enables a convenient means through which some of these challenges can be addressed. Here we describe an existing neural network approach [144] for state estimation which we modified to incorporate an external influence. This enables a hybrid estimator that can learn x_k from a mix of physiological data and an external influence. We show how this this method can be utilized when the features of interest are in the form of MPP observation pairs (m_k, r_k) and a continuous-valued observation s_k.

5.2 Data

We used the Neurological Status Assessment Dataset and the CFS and FMS Dataset for evaluating the state-space machine learning approach. Descriptions of these two datasets can be found in chapter 2 and chapter 4 respectively.

5.3 Methods

In [144], Krishnan et al., introduced a technique to learn x_k from data using neural networks. They consider the following general Gaussian state-space model:

$$x_k \sim \mathcal{N}(f_{\mu_x}(x_{k-1}), f_{\sigma_x^2}(x_{k-1})) \tag{138}$$

and

$$y_k \sim \Pi(f_y(x_k)). \tag{139}$$

Here y_k represents the observations. Both the state transition equation and the output equation are learned using two separate neural networks (for simplicity, we group both of them together under the title state-space neural network – SSNN). A separate recurrent neural network (RNN) is used to estimate x_k. Taking ψ and ϕ to denote the weights of the SSNN and the RNN respectively, the networks are trained by maximizing

$$\begin{aligned}\tilde{Q} = &\sum_{k=1}^{K} \mathbb{E}_{q_\phi(x_k|\vec{y})}\big[\log p_\psi(y_k|x_k)\big] - \mathrm{KL}(q_\phi(x_1|\vec{y})||p_\psi(x_1)) \\ &- \sum_{k=2}^{K} \mathbb{E}_{q_\phi(x_{k-1}|\vec{y})}\big[\mathrm{KL}(q_\phi(x_k|x_{k-1},\vec{y})||p_\psi(x_k|x_{k-1}))\big],\end{aligned} \tag{140}$$

where $p_\psi(\cdot)$ and $q_\phi(\cdot)$ denote density functions. The actual training is performed within the algorithm as a *minimization* of the negative term which we label Q. Analogous to a typical state-space method, with this alternate neural network-based approach, the SSNN replaces

the explicit state-space model, the RNN replaces the Bayesian filter and the weights of the neural networks replace the model parameters. Since neural networks are used to learn the state-space model, more complicated state transitions and input-output relationships are permitted. However, these come at the expense of interpretability.

The neural networks parameterized by ψ and ϕ are updated using stochastic backpropagation and details of this can be found in [144]. Here we point out some similarities with our regular state-space EM approach, particularly with regard to the terms in Q. For instance, when a binary variable m_k is present among the observations y_k, Q contains the summation

$$-\sum\left[m_k \log\left(\frac{1}{1+e^{-f_m(x_k)}}\right)+(1-m_k)\log\left(1-\frac{1}{1+e^{-f_m(x_k)}}\right)\right], \tag{141}$$

similar to the term observed in the log-likelihood of the regular EM algorithm, except now that $f_m(\cdot)$ is a function learned by the SSNN. Similarly, if a continuous-valued variable s_k is present in y_k, there is the summation

$$\sum \frac{1}{2}\log\left[2\pi f_{\sigma_s^2}(x_k)\right]+\frac{\left[s_k-f_{\mu_s}(x_k)\right]^2}{2f_{\sigma_s^2}(x_k)} \tag{142}$$

in Q, where $f_{\mu_s}(\cdot)$ and $f_{\sigma_s^2}(\cdot)$ represent mean and variance functions learned by the SSNN. Again, we have a similar term involving s_k in the log-likelihood of the EM algorithm.

We make two modifications to the original method in [144] here. We firstly change the training cost function to

$$\tilde{Q}=(1-\rho)Q+\rho\sum(x_k-l_k)^2, \tag{143}$$

where l_k is an external influence and $0\leq\rho\leq 1$. Our general approach is to first train the neural networks with $\rho=0$ similar to the original method in [144], then take the pre-trained model, change the value of ρ, introduce l_k and re-train for several more epochs to modify the state estimates. As we show later on, this permits a class of hybrid estimators to be

learned that can determine x_k from a mix of physiological data and an external influence. We also change the form of eq. (142) to sum only over the indices at which a point process event occur in the case of the MPP amplitudes r_k.

For the cortisol data, we use a neural network configuration for the SSNN with two layers and a hidden layer dimensionality of 256. We also selected 256 for the RNN size. This particular configuration was chosen because it yielded non-inverted state estimates. The moderately larger network size also is able to adapt to l_k quicker. On cortisol data, we use 500-epoch pre-trained neural networks with $\rho = 0$ and l_k is introduced thereafter.

For skin conductance, we selected a hidden layer dimensionality of 200 for the SSNN and two layers. The RNN size was also 200. We also use 2500-epoch pre-trained models prior to introducing l_k.

5.4 Results and Discussion

5.4.1 Cortisol Data

Cortisol-related energy state estimation results for CFS Patient 1 with and without (a hypothetical) l_k are shown in Fig. 22. Within each sub-figure, the sub-panels respectively depict: (a) the raw blood cortisol data z_k; (b) the deconvolved cortisol pulses (blue) and the estimated fit to r_k (red); (c) the reconstructed blood cortisol profile s_k (orange) and its estimated fit (red); (d) the probability of pulse occurrence p_k; (e) the energy production state x_k (purple) with its 95% confidence limits and the l_k term (blue). As illustrated in the figure, x_k can made to conform to l_k as closely as desired. Here we have chosen a binary-like l_k term corresponding to the sleep-wake times for the subjects merely for illustration [132]. We first trained the neural networks for 500 epochs with $\rho = 0$. We then took the pre-trained model, set $\rho = 0.5$, introduced l_k and trained for several more epochs. With $\rho = 0$, x_k is only estimated from the data. However, after setting $\rho = 0.5$ and training for an additional $j = 100$ epochs, x_k completely fits to the l_k. By changing ρ and the number of additional training epochs j, the l_k term can be allowed to have as much influence as

desired on the state estimates.

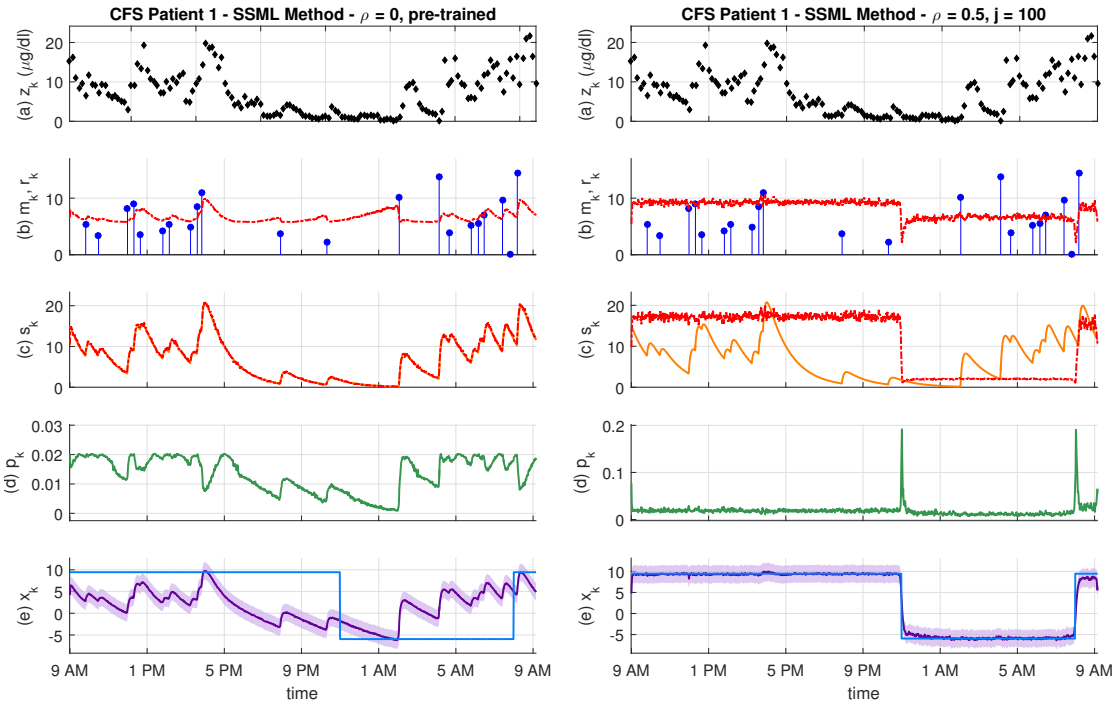

Figure 22: Effect of l_k on estimation when using state-space machine learning method.

A more practical example of the use of l_k in reducing overfitting during state estimation is shown in Fig. 23. Within each sub-figure, the sub-panels respectively depict: (a) the raw blood cortisol data z_k; (b) the deconvolved cortisol pulses (blue) and the estimated fit to r_k (red); (c) the reconstructed blood cortisol profile s_k (orange) and its estimated fit (red); (d) the probability of pulse occurrence p_k; (e) the energy production state x_k (purple) with its 95% confidence limits and the circadian l_k term (blue). For each subject, we first fit a circadian l_k term to the x_k estimate from the pre-trained model based on a least-squares minimization. Thereafter, we take the pre-trained model and train for an additional $j = 20$ epochs with $\rho = 0.75$. As can be seen from the sub-figures, the use of l_k helps undo some of the overfitting to s_k.

We also investigated the effect of the RNN size, the number of layers in the SSNN and

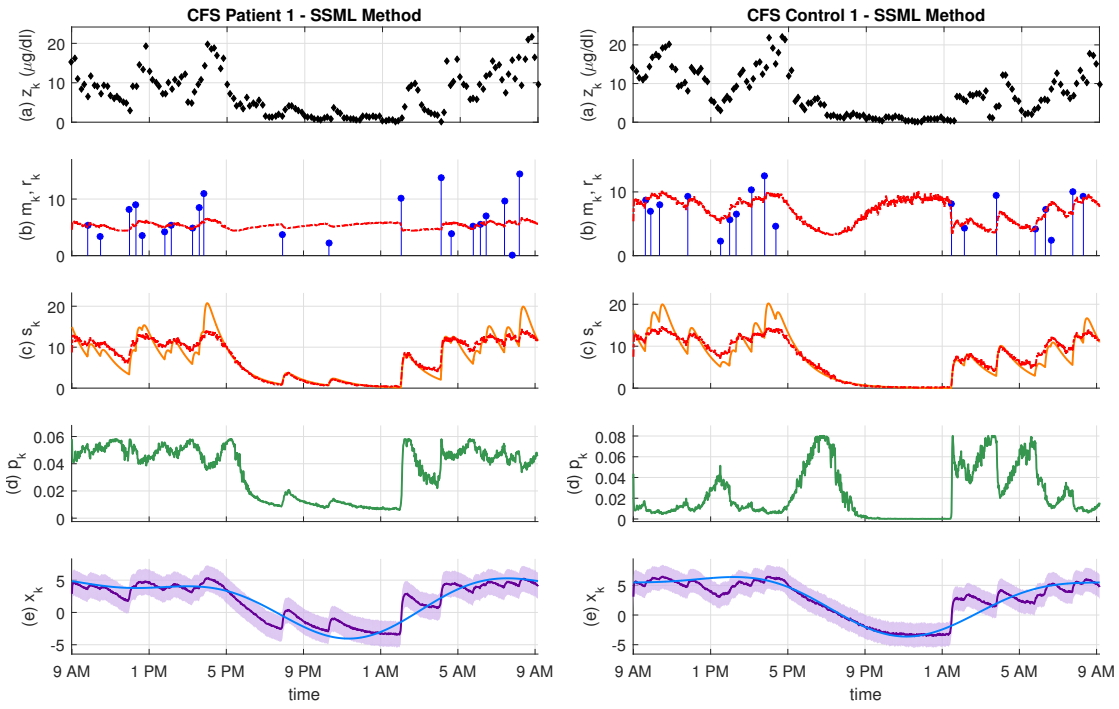

Figure 23: Cortisol-related energy state estimation using state-space machine learning method.

the dimensionality of the hidden layers in the SSNN on the resulting estimates (the number of layers in the two neural networks in the SSNN can be changed independently, but here we set them equal to each other and only changed them together). We did so using cortisol data from four subjects (CFS patient 1, CFS control 1, FMS patient 1 and FMS control 1). The following are some general observations we made:

- In a number of configurations, the neural networks learn a state x_k that is inverted. Physiologically, we would expect the cortisol-related energy production state to be high towards morning awakening and lower at bedtime. Even though inverted, the estimated x_k in these cases still, however, does tend to correspond to the general *shape* of the observed data. A similar inversion can also happen to the probability of pulse occurrence p_k. The neural networks can also learn an inverted state x_k with a

non-inverted p_k and vice versa.

- Whether a particular configuration with a certain RNN size and SSNN hidden layer dimensionality learns an inverted x_k or not tends to remain stable over subjects. For instance, if a neural network configuration with a hidden layer dimensionality of 256 in the SSNN and an RNN size of 256 learns a non-inverted state x_k for a particular subject, then that same configuration tends to provide non-inverted state estimates for the other three subjects as well. However, this is not always the case. Occasionally, a neural network configuration that provides a non-inverted state estimate for a particular subject will provide an inverted estimate for another subject. We examined this effect for 49×2 network configurations prior to selecting a suitable one (49 different combinations of the RNN size and SSNN hidden layer dimensionality for the cases of one layer and two layers in the SSNN). Manually exploring which of these configurations provide inverted or non-inverted state estimates is time consuming. As such, this trial-and-error approach would not be a suitable method for selecting a particular neural network configuration in the real-world. One possible alternative to this would be use an additional processing layer (perhaps another neural network) before the two main state-space neural network systems and configure this layer to handle the selection of the neural network sizes (layers and dimensionality) in an automated fashion.

- In general, a larger SSNN hidden layer dimensionality and RNN size tend to provide smoother estimates quicker (i.e., after a fewer number of training epochs). While the smaller network configurations can also provide similar smooth estimates, they tend to require more training epochs to do so. The larger neural network configurations are also able to adapt to the l_k term quicker. The larger networks do, however, tend to take more time to train.

- If the number of training epochs is kept fixed (e.g., 750), and we keep increasing the

number of *layers* in the SSNN (without changing number of *dimensionality* of the hidden layers), there comes a point where the fits to the observations (e.g., to s_k and p_k) become very noisy and continue to be so.

We chose the neural configurations described above for cortisol-related energy estimation and sympathetic arousal estimation taking into consideration some of the observations.

5.4.2 Skin Conductance Data

Sympathetic arousal estimates using the machine learning method for a representative subject in the Neurological Status Assessment Dataset are shown in Fig. 24. Here too we also illustrated two cases with and without l_k. Within each sub-figure, the sub-panels respectively depict: (a) the skin conductance signal z_k; (b) the deconvolved neural impulses (blue) and the estimated fit to r_k (red); (c) the tonic component s_k (orange) and its estimated fit (red); (d) the probability of impulse occurrence p_k; (e) the arousal state x_k (purple) with its 95% confidence limits (the l_k term is shown in blue on the right sub-figure); (f) HAI (the regions above 90% and below 10% are shaded in red and green respectively). The background colors correspond to the instruction period (red), the counting task (green), the Stroop test (orange), relaxation (blue) and emotional stress (yellow) respectively. As shown in the figure, the estimated arousal state can be made to conform to an external influence l_k as closely as we please in the case of skin conductance as well.

We evaluated the method for six participants in the dataset. The arousal estimates vary between different participants based on their individual responses to the stress stimuli. In general, for most of the participants, the sympathetic arousal estimates are high during the initial part of the experiment (mainly comprising of the cognitive stressors) and gradually diminish thereafter. A slight increase is to be noted around the time of the emotional stressor (horror movie clip). Variations between participants are also to be noted. The only exception to the general trend is participant 4. Unusually large increases were seen in the skin conductance signal for participant 4 at transition points in the experiment. These

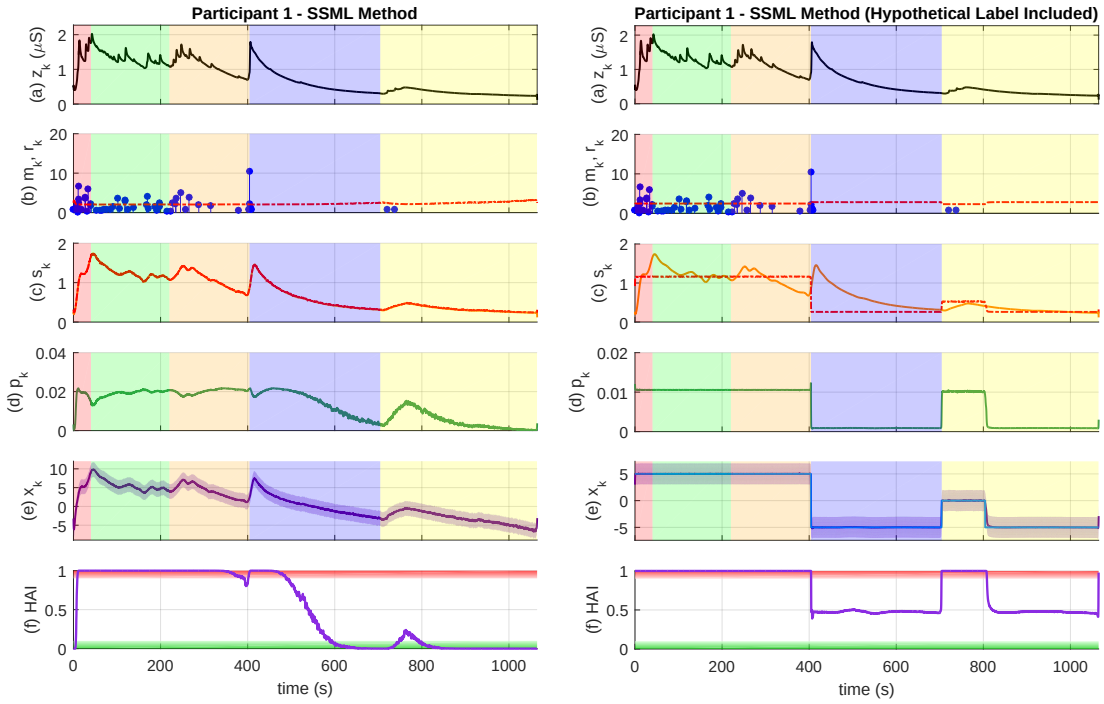

Figure 24: Sympathetic arousal estimation using state-space machine learning method.

unusual increases, possibly due to a form of motion artifact contamination, may have been the likely reason why the estimates deviated considerably from that of the others.

The cognitive stressors employed in the experiment (a counting task and a color-word association task) involve active engagement on the part of the subjects and likely is the reason for the elevated arousal during this portion. Watching the horror movie clip as a part of the emotional stressor only involves passive engagement and may not have generated an arousal response as large or prolonged.

Changes in internal states within the human body give rise to measurable electrical and chemical phenomena. It is likely, however, that physiological sensor data alone does not capture all of the necessary information required for estimating these latent states. Rather, behavioral measures (e.g., a person's social media posts, internet searches), medical expertise (e.g., knowledge of biological rhythms) etc. also remain a vital source of information

regarding what is occurring within the body. Consequently, physiological state estimation may ultimately need to be based on a fusion of sensor data and external information. The results shown in this chapter illustrate how latent states tied to markers in skin conductance and cortisol can be estimated based on physiological data and an external influence. The approaches could also find applications with other pulsatile phenomena as well.

The parameter ρ may be chosen based on the type of real-world application. One possible method of tuning ρ is to first test out the hybrid estimator on unseen data. Based on its performance on this data, an accuracy metric may be defined by how closely the state estimates agree with l_k. This accuracy metric can then be used to tune either ρ or any of the other hyper-parameters. Moreover, the general formulation

$$\tilde{Q} = (1-\rho)Q + \rho \sum (x_k - l_k)^2, \tag{144}$$

is not just applicable to the neural network approach but to the general state estimation framework as well. For instance, in all the other EM approaches we use, we have a similar log-likelihood term that is to be maximized. We typically isolate different components in this term, take derivatives and solve for the M-step updates. However, we could also obtain the M-step updates via an optimization of the full log-likelihood term. Thus, a quantity similar to $\rho \sum (x_k - l_k)^2$ could be added to this term as well. This would then enable the regular EM approach to also estimate latent states within the human body based on a mix of physiological data and an external influence.

6 Conclusions

Healthcare technology has seen rapid advances over the course of the past several years. In particular, a number of advances have been made in making healthcare systems smarter. It is likely that these smart technologies will play an increasing role in our everyday lives in the future. Wearable monitors are expected to play a crucial role here. In this dissertation, we argue that the basic system comprising of wearable sensing, diagnosis/detection and the prescription of corrective action forms a closed-loop framework. As such, tools from controls theory can be readily used to help build some of the smart healthcare systems of the future. Here we presented several state-space methods for estimating underlying states within the body that give rise to point process phenomena. These state estimators could be embedded within closed-loop controllers for automated therapy delivery in different scenarios.

6.1 Sympathetic Arousal and Skin Conductance

In our first section, we developed several methods for estimating sympathetic arousal from skin conductance measurements. These methods could be used to monitor patients with certain types of neuropsychiatric disorders (e.g., PTSD, depression) and perhaps even adjust closed-loop neurostimulation parameters accordingly. The methods could also find general stress management applications. The method based on treating the neural impulses underlying skin conductance variations as an MPP outperformed two of the other methods. We showed that the MPP-based approach did not run into the overfitting problem as did one of the other methods nor overreact to small impulses. Results were presented on two datasets having different stress conditions in experiments conducted both inside and outside laboratory environments. Our main contribution in this section was a set of point process state-space models and decoders for estimating sympathetic arousal from skin conductance. The models and decoders could have applicability elsewhere as well.

6.2 Sympathetic Arousal, Skin Conductance and Heart Rate

In the next section, we supplemented skin conductance features with heart rate measurements to estimate arousal from multi-timescale point process observations. Thus the primary research contribution involved the development of multi-timescale point process state-space models and decoders for tracking arousal. The methods help lay the groundwork for how other point process observations on different timescales may be incorporated into physiological state-space models. Respiration, for instance, is a slower point process but could very well be included for estimating sympathetic arousal as well. Results were shown on two different experimental datasets. The arousal estimates matched general expectations in the fear conditioning dataset and also revealed the interesting fact that task unfamiliarity may generate more stress than other factors (interruptions and time pressure) in the other dataset. The results overall indicate the viability of combining multi-timescale skin conductance and heart rate features within a single model for estimating arousal. Heart rate can be easily measured via PPG waveforms from the wrist. Skin conductance signals can also be recorded from this location. Thus a convenient wrist-worn wearable device could easily track arousal based on the models we developed.

6.3 Cortisol-related Energy Production

The hormone cortisol functions to increase blood glucose levels in response to external stressors. Cortisol therefore plays a key role in energy production within the body. The secretion dynamics of cortisol are very similar to that of skin conductance. Cortisol is secreted in pulses, just like neural impulses are responsible for skin conductance variations. We developed state-space methods for estimating an unobserved energy state from binary and continuous-valued blood cortisol measurements. When applied to data from patients with Cushing's disease, the results helped shed light on why hypercortisolism patients may experience daytime fatigue and nighttime sleeping difficulties. No significant differences were found in the case of CFS and FMS patients for whom there is also a lack of consensus

in the literature regarding cortisol differences. These state estimators could be embedded within automated drug infusion pumps for therapy delivery. Currently, many cortisol disorders are treated with oral medication doses and may not be optimal in certain cases. Here again, our main research contribution was a set of point process state-space models and decoders for estimating energy production from blood cortisol measurements. The methods could have more widespread applicability related to other pulsatile endocrine secretions as well.

6.4 Machine Learning for State Estimation

In our final section, we used machine learning for state estimation and addressed a notable drawback with traditional state-space methods. Traditional methods are unable to incorporate external influences such as subject-provided labels or domain expertise during estimation. We showed how an existing neural network approach for state estimation could be modified so as to yield a hybrid estimator capable of estimating a latent state based on a mix of physiological data and an external influence. Results were shown with both skin conductance and cortisol data. The basic framework for hybrid estimation could be applicable in a number of scenarios. For instance, the methodology could be used to incorporate medical expertise, patient feedback, location information etc., during estimation. The weighted combination could also be tuned so that the external influence is permitted to affect the state estimates as much as desired. We also pointed out how the parameter governing the degree of influence could be selected and how the basic mathematical framework for hybrid estimation could be adapted to the regular EM case as well. In this section, our main contribution involved developing a hybrid estimation method for physiological marked point process data capable of permitting an external influence to affect the state estimates.

6.5 Future Work

While our methods are able to track sympathetic arousal well, human emotion, in reality, is a complex multi-dimensional phenomena [10]. Therefore, an accurate characterization of a person's emotions requires the estimation of a vector-valued state (containing both arousal and valence). Future work would include the extension of the current models to estimate valence alongside arousal. Moreover, here we have also only considered a single hormone for energy state estimation. Energy production and calorie intake are influenced by more than one hormone in the body. Therefore, the models presented here could be extended to include features from more hormones and perhaps even estimate a vector-valued energy state (e.g., with two components related to calorie intake and stress). Furthermore, the human body is also not static; instead changes occur over time. Consequently, a model trained on data today may not work accurately several weeks from now. The models described here are based on parameters estimated on fixed segments of data. Therefore, another possible future direction of research is to introduce time-varying parameters into the models. The results we have presented thus far have also largely been based on data acquired from healthy subjects or previously-collected data. We do envisage that the methods we have developed could be utilized for treating actual patients. These may include patients diagnosed with neuropsychiatric or hormone disorders. Real-world implementation of the methods could be pursued in future with appropriate clinical collaboration.

References

[1] A. B. Martin, M. Hartman, B. Washington, A. Catlin, and National Health Expenditure Accounts Team, "National health care spending in 2017: Growth slows to post–great recession rates; share of GDP stabilizes," *Health Affairs*, pp. 10–1377, 2018.

[2] A. M. Sisko, S. P. Keehan, J. A. Poisal, G. A. Cuckler, S. D. Smith, A. J. Madison, K. E. Rennie, and J. C. Hardesty, "National health expenditure projections, 2018–27: Economic and demographic trends drive spending and enrollment growth," *Health Affairs*, vol. 38, no. 3, pp. 491–501, 2019.

[3] OECD, *Health at a Glance 2019*, 2019. [Online]. Available: https://www.oecd-ilibrary.org/content/publication/4dd50c09-en

[4] S. Park and S. Jayaraman, "Wearable biomedical systems and mHealth," in *Telemedicine and Electronic Medicine*. CRC Press, 2018, pp. 134–159.

[5] A. S. Miller, J. A. Cafazzo, and E. Seto, "A game plan: Gamification design principles in mHealth applications for chronic disease management," *Health Informatics Journal*, vol. 22, no. 2, pp. 184–193, 2016.

[6] J. Wu, H. Li, S. Cheng, and Z. Lin, "The promising future of healthcare services: When big data analytics meets wearable technology," *Information & Management*, vol. 53, no. 8, pp. 1020–1033, 2016.

[7] J. Kim, A. S. Campbell, B. E.-F. de Ávila, and J. Wang, "Wearable biosensors for healthcare monitoring," *Nature Biotechnology*, vol. 37, no. 4, pp. 389–406, 2019.

[8] National Alliance on Mental Health, "Mental health by the numbers," https://www.nami.org/learn-more/mental-health-by-the-numbers, Retrieved: November 16, 2019.

[9] S. Moratti, G. Rubio, P. Campo, A. Keil, and T. Ortiz, "Hypofunction of right temporoparietal cortex during emotional arousal in depression," *Archives of General Psychiatry*, vol. 65, no. 5, pp. 532–541, 2008.

[10] J. A. Russell, "A circumplex model of affect," *Journal of Personality and Social Psychology*, vol. 39, no. 6, p. 1161, 1980.

[11] A. L. Mahan and K. J. Ressler, "Fear conditioning, synaptic plasticity and the amygdala: Implications for posttraumatic stress disorder," *Trends in Neurosciences*, vol. 35, no. 1, pp. 24–35, 2012.

[12] R. Belmaker, "Bipolar disorder," *New England Journal of Medicine*, vol. 351, no. 5, pp. 476–486, 2004.

[13] R. W. Picard, *Affective Computing*. MIT press, 2000.

[14] H. J. Pijeira-Díaz, H. Drachsler, S. Järvelä, and P. A. Kirschner, "Sympathetic arousal commonalities and arousal contagion during collaborative learning: How attuned are triad members?" *Computers in Human Behavior*, vol. 92, pp. 188–197, 2019.

[15] M.-Z. Poh, N. C. Swenson, and R. W. Picard, "A wearable sensor for unobtrusive, long-term assessment of electrodermal activity," *IEEE Transactions on Biomedical Engineering*, vol. 57, no. 5, pp. 1243–1252, 2010.

[16] D. Puigdemont, R. Pérez-Egea, M. J. Portella, J. Molet, J. de Diego-Adeliño, A. Gironell, J. Radua, B. Gómez-Anson, R. Rodríguez, M. Serra, C. de Quintana, F. Artigas, E. Álvarez, and V. Pérez, "Deep brain stimulation of the subcallosal cingulate gyrus: Further evidence in treatment-resistant major depression," *International Journal of Neuropsychopharmacology*, vol. 15, no. 1, pp. 121–133, 2012.

[17] A. Merkl, G.-H. Schneider, T. Schönecker, S. Aust, K.-P. Kühl, A. Kupsch, A. A. Kühn, and M. Bajbouj, "Antidepressant effects after short-term and chronic stimulation of the subgenual cingulate gyrus in treatment-resistant depression," *Experimental Neurology*, vol. 249, pp. 160–168, 2013.

[18] J.-P. Langevin, J. Chen, R. Koek, D. Sultzer, M. Mandelkern, H. Schwartz, and S. Krahl, "Deep brain stimulation of the basolateral amygdala: Targeting technique and electrodiagnostic findings," *Brain Sciences*, vol. 6, no. 3, p. 28, 2016.

[19] R. J. Koek, J.-P. Langevin, S. E. Krahl, H. J. Kosoyan, H. N. Schwartz, J. W. Chen, R. Melrose, M. J. Mandelkern, and D. Sultzer, "Deep brain stimulation of the basolateral amygdala for treatment-refractory combat post-traumatic stress disorder (PTSD): Study protocol for a pilot randomized controlled trial with blinded, staggered onset of stimulation," *Trials*, vol. 15, no. 1, p. 356, 2014.

[20] S. Little, A. Pogosyan, S. Neal, B. Zavala, L. Zrinzo, M. Hariz, T. Foltynie, P. Limousin, K. Ashkan, J. FitzGerald, A. L. Green, T. Z. Aziz, and P. Brown, "Adaptive deep brain stimulation in advanced Parkinson disease," *Annals of Neurology*, vol. 74, no. 3, pp. 449–457, 2013.

[21] S. Little, M. Beudel, L. Zrinzo, T. Foltynie, P. Limousin, M. Hariz, S. Neal, B. Cheeran, H. Cagnan, J. Gratwicke, T. Z. Aziz, A. Pogosyan, and P. Brown, "Bilateral adaptive deep brain stimulation is effective in Parkinson's disease," *Journal of Neurology, Neurosurgery & Psychiatry*, vol. 87, no. 7, pp. 717–721, 2016.

[22] B. Rosin, M. Slovik, R. Mitelman, M. Rivlin-Etzion, S. N. Haber, Z. Israel, E. Vaadia, and H. Bergman, "Closed-loop deep brain stimulation is superior in ameliorating Parkinsonism," *Neuron*, vol. 72, no. 2, pp. 370–384, 2011.

[23] H. D. Critchley, R. N. Melmed, E. Featherstone, C. J. Mathias, and R. J. Dolan, "Volitional control of autonomic arousal: A functional magnetic resonance study," *Neuroimage*, vol. 16, no. 4, pp. 909–919, 2002.

[24] D. S. Wickramasuriya, M. R. Amin, and R. T. Faghih, "Skin conductance as a viable alternative for closing the deep brain stimulation loop in neuropsychiatric disorders," *Frontiers in Neuroscience*, vol. 13, p. 780, 2019.

[25] A. Clow, F. Hucklebridge, and L. Thorn, "The cortisol awakening response in context," in *Science of Awakening*, ser. International Review of Neurobiology, A. Clow and L. Thorn, Eds. Academic Press, 2010, vol. 93, pp. 153 – 175. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0074774210930079

[26] E. K. Do Yup Lee and M. H. Choi, "Technical and clinical aspects of cortisol as a biochemical marker of chronic stress," *BMB Reports*, vol. 48, no. 4, p. 209, 2015.

[27] J. D. Veldhuis, "Pulsatile hormone secretion: mechanisms, significance and evaluation," in *Ultradian Rhythms from Molecules to Mind.* Springer, 2008, pp. 229–248.

[28] R. T. Faghih, M. A. Dahleh, G. K. Adler, E. B. Klerman, and E. N. Brown, "Deconvolution of serum cortisol levels by using compressed sensing," *PloS One*, vol. 9, no. 1, p. e85204, 2014.

[29] J. D. Veldhuis, A. Iranmanesh, G. Lizarralde, and M. L. Johnson, "Amplitude modulation of a burstlike mode of cortisol secretion subserves the circadian glucocorticoid rhythm," *American Journal of Physiology-Endocrinology and Metabolism*, vol. 257, no. 1, pp. E6–E14, 1989.

[30] E. N. Brown, P. M. Meehan, and A. P. Dempster, "A stochastic differential equation model of diurnal cortisol patterns," *American Journal of Physiology-Endocrinology and Metabolism*, vol. 280, no. 3, pp. E450–E461, 2001.

[31] R. T. Faghih, M. A. Dahleh, and E. N. Brown, "An optimization formulation for characterization of pulsatile cortisol secretion," *Frontiers in Neuroscience*, vol. 9, p. 228, 2015.

[32] A. A. Nella, A. Mallappa, A. F. Perritt, V. Gounden, P. Kumar, N. Sinaii, L.-A. Daley, A. Ling, C.-Y. Liu, S. J. Soldin, and D. P. Merke, "A phase 2 study of continuous subcutaneous hydrocortisone infusion in adults with congenital adrenal hyperplasia," *The Journal of Clinical Endocrinology & Metabolism*, vol. 101, no. 12, pp. 4690–4698, 2016.

[33] A. Khanna, R. Khurana, A. Kyriacou, R. Davies, and D. Ray, "Management of adrenocortical insufficiency with continuous subcutaneous hydrocortisone infusion: Long-term experience in three patients," *Endocrinology, Diabetes & Metabolism Case Reports*, vol. 2015, no. 1, 2015.

[34] A. Mallappa, A. A. Nella, N. Sinaii, H. Rao, V. Gounden, A. F. Perritt, P. Kumar, A. Ling, C.-Y. Liu, S. J. Soldin, and D. P. Merke, "Long-term use of continuous subcutaneous hydrocortisone infusion therapy in patients with congenital adrenal hyperplasia," *Clinical Endocrinology*, vol. 89, no. 4, pp. 399–407, 2018.

[35] US Food and Drug Administration, "Clinical decision support software," https://www.fda.gov/regulatory-information/search-fda-guidance-documents/clinical-decision-support-software, Retrieved: November 16, 2019.

[36] R. T. Faghih, P. A. Stokes, M.-F. Marin, R. G. Zsido, S. Zorowitz, B. L. Rosenbaum, H. Song, M. R. Milad, D. D. Dougherty, E. N. Eskandar, and R. Barbieri, "Characterization of fear conditioning and fear extinction by analysis of electrodermal activity," in *37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 2015, pp. 7814–7818.

[37] M. R. Amin and R. T. Faghih, "Sparse deconvolution of electrodermal activity via continuous-time system identification," *IEEE Transactions on Biomedical Engineering*, vol. 66, no. 9, pp. 2585–2595, 2019.

[38] M. Benedek and C. Kaernbach, "Decomposition of skin conductance data by means of nonnegative deconvolution," *Psychophysiology*, vol. 47, no. 4, pp. 647–658, 2010.

[39] A. C. Smith and E. N. Brown, "Estimating a state-space model from point process observations," *Neural computation*, vol. 15, no. 5, pp. 965–991, 2003.

[40] E. N. Brown, L. M. Frank, D. Tang, M. C. Quirk, and M. A. Wilson, "A statistical paradigm for neural spike train decoding applied to position prediction from ensemble

firing patterns of rat hippocampal place cells," *Journal of Neuroscience*, vol. 18, no. 18, pp. 7411–7425, 1998.

[41] A. C. Smith, L. M. Frank, S. Wirth, M. Yanike, D. Hu, Y. Kubota, A. M. Graybiel, W. A. Suzuki, and E. N. Brown, "Dynamic analysis of learning in behavioral experiments," *Journal of Neuroscience*, vol. 24, no. 2, pp. 447–461, 2004.

[42] A. C. Smith, M. R. Stefani, B. Moghaddam, and E. N. Brown, "Analysis and design of behavioral experiments to characterize population learning," *Journal of Neurophysiology*, vol. 93, no. 3, pp. 1776–1792, 2005.

[43] M. J. Prerau, A. C. Smith, U. T. Eden, M. Yanike, W. A. Suzuki, and E. N. Brown, "A mixed filter algorithm for cognitive state estimation from simultaneously recorded continuous and binary measures of performance," *Biological Cybernetics*, vol. 99, no. 1, pp. 1–14, 2008.

[44] M. J. Prerau, A. C. Smith, U. T. Eden, Y. Kubota, M. Yanike, W. Suzuki, A. M. Graybiel, and E. N. Brown, "Characterizing learning by simultaneous analysis of continuous and binary measures of performance," *Journal of Neurophysiology*, vol. 102, no. 5, pp. 3060–3072, 2009.

[45] T. P. Coleman, M. Yanike, W. Suzuki, and E. N. Brown, "A mixed filter algorithm for state estimation from simultaneously recorded continuous-valued point process and binary observations," in *40th Asilomar Conference on Signals, Systems and Computers*, 2006, pp. 1949–1953.

[46] T. P. Coleman, M. Yanike, W. A. Suzuki, and E. N. Brown, "A mixed-filter algorithm for dynamically tracking learning from multiple behavioral and neurophysiological measures," *The dynamic brain: An exploration of neuronal variability and its functional significance*, pp. 3–28, 2011.

[47] Y. Yang and M. M. Shanechi, "An adaptive brain-machine interface algorithm for control of burst suppression in medical coma," in *2014 36th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 2014, pp. 1638–1641.

[48] M. M. Shanechi, J. J. Chemali, M. Liberman, K. Solt, and E. N. Brown, "A brain-machine interface for control of burst suppression in medical coma," in *35th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 2013, pp. 1575–1578.

[49] M. M. Shanechi, J. J. Chemali, M. Liberman, K. Solt, and E. N. Brown, "A brain-machine interface for control of medically-induced coma," *PLoS Computational Biology*, vol. 9, no. 10, p. e1003284, 2013.

[50] Y. Yang and M. M. Shanechi, "An adaptive and generalizable closed-loop system for control of medically induced coma and other states of anesthesia," *Journal of Neural Engineering*, vol. 13, no. 6, p. 066019, 2016.

[51] Y. Yang, J. T. Lee, J. A. Guidera, K. Y. Vlasov, J. Pei, E. N. Brown, K. Solt, and M. M. Shanechi, "Developing a personalized closed-loop controller of medically-induced coma in a rodent model," *Journal of Neural Engineering*, vol. 16, no. 3, p. 036022, 2019.

[52] P. L. Purdon, E. T. Pierce, E. A. Mukamel, M. J. Prerau, J. L. Walsh, K. F. K. Wong, A. F. Salazar-Gomez, P. G. Harrell, A. L. Sampson, A. Cimenser, S. Ching, N. J. Kopell, C. Tavares-Stoeckel, K. Habeeb, R. Merhar, and E. N. Brown, "Electroencephalogram signatures of loss and recovery of consciousness from propofol," *Proceedings of the National Academy of Sciences*, vol. 110, no. 12, pp. E1142–E1151, 2013.

[53] K. F. K. Wong, A. C. Smith, E. T. Pierce, P. G. Harrell, J. L. Walsh, A. F. Salazar, C. L. Tavares, A. Cimenser, M. J. Prerau, E. A. Mukamel, A. Sampson, P. L. Purdon, and E. N. Brown, "Bayesian analysis of trinomial data in behavioral experiments and its application to human studies of general anesthesia," in *Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 2011, pp. 4705–4708.

[54] Y. Yang and M. M. Shanechi, "A generalizable adaptive brain-machine interface design for control of anesthesia," in *37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 2015, pp. 1099–1102.

[55] M. M. Shanechi, Z. M. Williams, G. W. Wornell, R. C. Hu, M. Powers, and E. N. Brown, "A real-time brain-machine interface combining motor target and trajectory intent using an optimal feedback control design," *PloS One*, vol. 8, no. 4, p. e59049, 2013.

[56] M. M. Shanechi, R. C. Hu, M. Powers, G. W. Wornell, E. N. Brown, and Z. M. Williams, "Neural population partitioning and a concurrent brain-machine interface for sequential motor function," *Nature Neuroscience*, vol. 15, no. 12, p. 1715, 2012.

[57] M. M. Shanechi, A. L. Orsborn, and J. M. Carmena, "Robust brain-machine interface design using optimal feedback control modeling and adaptive point process filtering," *PLoS computational biology*, vol. 12, no. 4, p. e1004730, 2016.

[58] M. M. Shanechi, G. W. Wornell, Z. M. Williams, and E. N. Brown, "Feedback-controlled parallel point process filter for estimation of goal-directed movements from neural signals," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 21, no. 1, pp. 129–140, 2012.

[59] M. J. Prerau, K. E. Hartnack, G. Obregon-Henao, A. Sampson, M. Merlino, K. Gannon, M. T. Bianchi, J. M. Ellenbogen, and P. L. Purdon, "Tracking the sleep onset process: An empirical model of behavioral and physiological dynamics," *PLoS Computational Biology*, vol. 10, no. 10, p. e1003866, 2014.

[60] P. A. Low, "Chapter 51 - Sweating," in *Primer on the Autonomic Nervous System (Third Edition)*, D. Robertson, I. Biaggioni, G. Burnstock, P. A. Low, and J. F. Paton, Eds. San Diego: Academic Press, 2012, pp. 249 – 251.

[61] H. D. Critchley, "Electrodermal responses: What happens in the brain," *The Neuroscientist*, vol. 8, no. 2, pp. 132–142, 2002.

[62] D. E. Aikins, D. C. Johnson, J. L. Borelli, D. H. Klemanski, P. M. Morrissey, T. L. Benham, S. M. Southwick, and D. F. Tolin, "Thought suppression failures in combat PTSD: A cognitive load hypothesis," *Behaviour Research and Therapy*, vol. 47, no. 9, pp. 744–751, 2009.

[63] J. C. Laberg and B. Ellertsen, "Psychophysiological indicators of craving in alcoholics: Effects of cue exposure," *British Journal of Addiction*, vol. 82, no. 12, pp. 1341–1348, 1987.

[64] K. Kallinen and N. Ravaja, "Emotion-related effects of speech rate and rising vs. falling background music melody during audio news: The moderating influence of personality," *Personality and Individual Differences*, vol. 37, no. 2, pp. 275–288, 2004.

[65] C. Lithari, C. Frantzidis, C. Papadelis, A. B. Vivas, M. Klados, C. Kourtidou-Papadeli, C. Pappas, A. Ioannides, and P. Bamidis, "Are females more responsive to emotional stimuli? A neurophysiological study across arousal and valence dimensions," *Brain Topography*, vol. 23, no. 1, pp. 27–40, 2010.

[66] N. Mella, L. Conty, and V. Pouthas, "The role of physiological arousal in time perception: Psychophysiological evidence from an emotion regulation paradigm," *Brain and Cognition*, vol. 75, no. 2, pp. 182–187, 2011.

[67] Y. Nagai, H. D. Critchley, E. Featherstone, M. R. Trimble, and R. J. Dolan, "Activity in ventromedial prefrontal cortex covaries with sympathetic skin conductance level: A physiological account of a "default mode" of brain function," *Neuroimage*, vol. 22, no. 1, pp. 243–251, 2004.

[68] L. M. Gatzke-Kopp, A. Raine, R. Loeber, M. Stouthamer-Loeber, and S. R. Steinhauer, "Serious delinquent behavior, sensation seeking, and electrodermal arousal," *Journal of Abnormal Child Psychology*, vol. 30, no. 5, pp. 477–486, 2002.

[69] R. J. Barry and E. N. Sokolov, "Habituation of phasic and tonic components of the orienting reflex," *International Journal of Psychophysiology*, vol. 15, no. 1, pp. 39–42, 1993.

[70] S. D. Kreibig, "Autonomic nervous system activity in emotion: A review," *Biological Psychology*, vol. 84, no. 3, pp. 394–421, 2010.

[71] D. S. Wickramasuriya, C. Qi, and R. T. Faghih, "A state-space approach for detecting stress from electrodermal activity," in *2018 40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, July 2018, pp. 3562–3567.

[72] D. S. Wickramasuriya and R. T. Faghih, "A Bayesian filtering approach for tracking arousal from binary and continuous skin conductance features," *IEEE Transactions on Biomedical Engineering*, vol. 67, no. 6, pp. 1749–1760, 2020.

[73] D. S. Wickramasuriya and R. T. Faghih, "A marked point process filtering approach for tracking sympathetic arousal from skin conductance," *IEEE Access*, vol. 8, pp. 68 499–68 513, 2020.

[74] J. Birjandtalab, D. Cogan, M. B. Pouyan, and M. Nourani, "A non-EEG biosignals dataset for assessment and visualization of neurological status," in *2016 IEEE International Workshop on Signal Processing Systems (SiPS)*. IEEE, 2016, pp. 110–114.

[75] J. Tulen, P. Moleman, H. Van Steenis, and F. Boomsma, "Characterization of stress reactions to the Stroop color word test," *Pharmacology Biochemistry and Behavior*, vol. 32, no. 1, pp. 9–15, 1989.

[76] J. A. Healey and R. W. Picard, "Detecting stress during real-world driving tasks using physiological sensors," *IEEE Transactions on Intelligent Transportation Systems*, vol. 6, no. 2, pp. 156–166, 2005.

[77] A. Greco, G. Valenza, A. Lanata, E. P. Scilingo, and L. Citi, "cvxEDA: A convex optimization approach to electrodermal activity processing," *IEEE Transactions on Biomedical Engineering*, vol. 63, no. 4, pp. 797–804, 2016.

[78] F. Ershad, A. Thukral, J. Yue, P. Comeaux, Y. Lu, H. Shim, K. Sim, N.-I. Kim, Z. Rao, R. Guevara, L. Contreras, F. Pan, Y. Zhang, Y.-S. Guan, P. Yang, X. Wang, P. Wang, X. Wu, and C. Yu, "Ultra-conformal drawn-on-skin electronics for multifunctional motion artifact-free sensing and point-of-care treatment," *Nature Communications*, vol. 11, no. 1, pp. 1–13, 2020.

[79] X. Deng, R. T. Faghih, R. Barbieri, A. C. Paulk, W. F. Asaad, E. N. Brown, D. D. Dougherty, A. S. Widge, E. N. Eskandar, and U. T. Eden, "Estimating a dynamic state to relate neural spiking activity to behavioral signals during cognitive tasks," in *37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 2015, pp. 7808–7813.

[80] P. McCullagh and J. A. Nelder, *Generalized Linear Models*. CRC press, 1989, vol. 37.

[81] J. J. Braithwaite, D. G. Watson, R. Jones, and M. Rowe, "A guide for analysing electrodermal activity (EDA) & skin conductance responses (SCRs) for psychological experiments," *Psychophysiology*, vol. 49, no. 1, pp. 1017–1034, 2013.

[82] W. Boucsein, *Electrodermal Activity*. Springer Science & Business Media, 2012.

[83] D. R. Bach, G. Flandin, K. J. Friston, and R. J. Dolan, "Modelling event-related skin conductance responses," *International Journal of Psychophysiology*, vol. 75, no. 3, pp. 349–356, 2010.

[84] M. E. Dawson, A. M. Schell, and D. L. Filion, "The electrodermal system," *Handbook of Psychophysiology*, vol. 2, pp. 200–223, 2007.

[85] J. M. Mendel, *Lessons in Estimation Theory for Signal Processing, Communications and Control*. Pearson Education, 1995.

[86] P. D. Jong and M. J. Mackinnon, "Covariances for smoothed estimates in state-space models," *Biometrika*, vol. 75, no. 3, pp. 601–602, 1988.

[87] Y. Nishiyama, M. Kanagawa, A. Gretton, and K. Fukumizu, "Model-based kernel sum rule: Kernel Bayesian inference with probabilistic models," *Machine Learning*, pp. 1–34, 2020.

[88] Y. Nishiyama, A. Afsharinejad, S. Naruse, B. Boots, and L. Song, "The non-parametric kernel Bayes smoother," in *Artificial Intelligence and Statistics*, 2016, pp. 547–555.

[89] M. Kanagawa, Y. Nishiyama, A. Gretton, and K. Fukumizu, "Filtering with state-observation examples via kernel Monte Carlo filter," *Neural Computation*, vol. 28, no. 2, pp. 382–444, 2016.

[90] K. Fukumizu, L. Song, and A. Gretton, "Kernel Bayes' rule: Bayesian inference with positive definite kernels," *The Journal of Machine Learning Research*, vol. 14, no. 1, pp. 3753–3783, 2013.

[91] J. E. Sharman, A. P. Avolio, J. Baulmann, A. Benetos, J. Blacher, C. L. Blizzard, P. Boutouyrie, C.-H. Chen, P. Chowienczyk, J. R. Cockcroft, J. K. Cruickshank, I. Ferreira, L. Ghiadoni, A. Hughes, P. Jankowski, S. Laurent, B. J. McDonnell, C. McEniery, S. C. Millasseau, T. G. Papaioannou, G. Parati, J. B. Park, A. D. Protogerou, M. J. Roman, G. Schillaci, P. Segers, G. S. Stergiou, H. Tomiyama, R. R. Townsend, L. M. Van Bortel, J. Wang, S. Wassertheurer, T. Weber, I. B. Wilkinson, and C. Vlachopoulos, "Validation of non-invasive central blood pressure devices: ARTERY society task force consensus statement on protocol standardization," *European Heart Journal*, vol. 38, no. 37, pp. 2805–2812, 2017.

[92] J. E. Hall, *Guyton and Hall Textbook of Medical Physiology*. Elsevier Health Sciences, 2016.

[93] D. S. Wickramasuriya and R. T. Faghih, "A novel filter for tracking real-world cognitive stress using multi-time-scale point process observations," in *41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, July 2019, pp. 599–602.

[94] D. S. Wickramasuriya and R. T. Faghih, "A mixed filter algorithm for sympathetic arousal tracking from skin conductance and heart rate measurements in Pavlovian fear conditioning," *Plos One*, vol. 15, no. 4, p. e0231659, 2020.

[95] R. C. Drew and L. I. Sinoway, "Autonomic control of the heart," in *Primer on the Autonomic Nervous System*. Elsevier, 2012, pp. 177–180.

[96] G. G. Berntson, J. T. Cacioppo, and K. S. Quigley, "The metrics of cardiac

chronotropism: Biometric perspectives," *Psychophysiology*, vol. 32, no. 2, pp. 162–171, 1995.

[97] G. G. Berntson, J. Thomas Bigger Jr, D. L. Eckberg, P. Grossman, P. G. Kaufmann, M. Malik, H. N. Nagaraja, S. W. Porges, J. P. Saul, P. H. Stone, and M. W. Van der Molen, "Heart rate variability: Origins, methods, and interpretive caveats," *Psychophysiology*, vol. 34, no. 6, pp. 623–648, 1997.

[98] R. Barbieri, E. C. Matten, A. A. Alabi, and E. N. Brown, "A point-process model of human heartbeat intervals: New definitions of heart rate and heart rate variability," *American Journal of Physiology-Heart and Circulatory Physiology*, vol. 288, no. 1, pp. H424–H435, 2005.

[99] G. B. Stanley, K. Poolla, and R. A. Siegel, "Threshold modeling of autonomic control of heart rate variability," *IEEE Transactions on Biomedical Engineering*, vol. 47, no. 9, pp. 1147–1153, 2000.

[100] R. Barbieri and E. N. Brown, "Analysis of heartbeat dynamics by point process adaptive filtering," *IEEE Transactions on Biomedical Engineering*, vol. 53, no. 1, pp. 4–12, 2006.

[101] A. Boardman, F. S. Schlindwein, and A. P. Rocha, "A study on the optimum order of autoregressive models for heart rate variability," *Physiological Measurement*, vol. 23, no. 2, p. 325, 2002.

[102] A. Pichon, M. Roulaud, S. Antoine-Jonville, C. de Bisschop, and A. Denjean, "Spectral analysis of heart rate variability: Interchangeability between autoregressive analysis and fast Fourier transform," *Journal of Electrocardiology*, vol. 39, no. 1, pp. 31–37, 2006.

[103] R. Barbieri and E. N. Brown, "Application of dynamic point process models to cardiovascular control," *Biosystems*, vol. 93, no. 1-2, pp. 120–125, 2008.

[104] U. T. Eden, L. Srinivasan, and S. V. Sarma, "Nueral signal processing tutorial II: Point process model estimation and goodness-of-fit analysis," in *Neural Signal Processing: Quantitative Analysis of Neural Activity*, P. Mitra, Ed. Washington DC: Society for Neuroscience, 2008, ch. 9, pp. 79–87.

[105] S. Koldijk, M. Sappelli, S. Verberne, M. A. Neerincx, and W. Kraaij, "The SWELL knowledge work dataset for stress and user modeling research," in *16th International Conference on Multimodal Interaction*. ACM, 2014, pp. 291–298.

[106] S. Doberenz, W. T. Roth, E. Wollburg, N. I. Maslowski, and S. Kim, "Methodological considerations in ambulatory skin conductance monitoring," *International Journal of Psychophysiology*, vol. 80, no. 2, pp. 87–95, 2011.

[107] S. Maren, "Neurobiology of Pavlovian fear conditioning," *Annual Review of Neuroscience*, vol. 24, no. 1, pp. 897–931, 2001.

[108] M. R. Milad, S. Igoe, and S. P. Orr, "Fear conditioning in rodents and humans," in *Animal Models of Behavioral Analysis*. Springer, 2011, pp. 111–132.

[109] O. V. Lipp, "Human fear learning: Contemporary procedures and measurement," *Fear and Learning: From Basic Processes to Clinical Implications*, no. 2001, pp. 37–51, 2006.

[110] A. Tzovara, N. Hofer, D. R. Bach, G. Castegnetti, S. Gerster, C. W. Korn, P. C. Paulus, and M. Staib, "PsPM-TC: SCR, ECG, EMG and respiration measurements in a discriminant trace fear conditioning task with visual CS and electrical US." Aug. 2018. [Online]. Available: https://doi.org/10.5281/zenodo.1404810

[111] G. Castegnetti, A. Tzovara, M. Staib, P. C. Paulus, N. Hofer, and D. R. Bach, "Modeling fear-conditioned bradycardia in humans," *Psychophysiology*, vol. 53, no. 6, pp. 930–939, 2016.

[112] G. Castegnetti, A. Tzovara, M. Staib, S. Gerster, and D. R. Bach, "Assessing fear learning via conditioned respiratory amplitude responses," *Psychophysiology*, vol. 54, no. 2, pp. 215–223, 2017.

[113] A. Tzovara, C. W. Korn, and D. R. Bach, "Human Pavlovian fear conditioning conforms to probabilistic learning," *PLoS Computational Biology*, vol. 14, no. 8, p. e1006243, 2018.

[114] E. N. Brown, R. Barbieri, V. Ventura, R. E. Kass, and L. M. Frank, "The time-rescaling theorem and its application to neural spike train data analysis," *Neural Computation*, vol. 14, no. 2, pp. 325–346, 2002.

[115] U. T. Eden, L. M. Frank, R. Barbieri, V. Solo, and E. N. Brown, "Dynamic analysis of neural encoding by point process adaptive filtering," *Neural Computation*, vol. 16,

no. 5, pp. 971–998, 2004.

[116] R. Barbieri, M. C. Quirk, L. M. Frank, M. A. Wilson, and E. N. Brown, "Construction and analysis of non-Poisson stimulus-response models of neural spiking activity," *Journal of Neuroscience Methods*, vol. 105, no. 1, pp. 25–37, 2001.

[117] S. Koyama and R. E. Kass, "Spike train probability models for stimulus-driven leaky integrate-and-fire neurons," *Neural Computation*, vol. 20, no. 7, pp. 1776–1795, 2008.

[118] L. B. Baker, "Physiology of sweat gland function: The roles of sweating and sweat composition in human health," *Temperature*, vol. 6, no. 3, pp. 211–259, 2019.

[119] H. Kotses and K. D. Glaus, "Latency of multiple skin conductance responses in differential classical conditioning," *Biological Psychology*, vol. 5, no. 1, pp. 1–6, 1977.

[120] J. E. Koepke and K. H. Pribram, "Habituation of GSR as a function of stimulus duration and spontaneous activity." *Journal of Comparative and Physiological Psychology*, vol. 61 3, pp. 442–448, 1966.

[121] M. C. Morrow, "Delayed CR in classical GSR conditioning," *Psychological Reports*, vol. 29, no. 3, pp. 1003–1009, 1971.

[122] A. Esposito and V. Bianchi, *Cortisol: Physiology, Regulation and Health Implications*. Nova Science Publishers, 2012.

[123] R. M. Torrente-Rodríguez, J. Tu, Y. Yang, J. Min, M. Wang, Y. Song, Y. Yu, C. Xu, C. Ye, W. W. IsHak, and W. Gao, "Investigation of cortisol dynamics in human sweat using a graphene-based wireless mHealth system," *Matter*, 2020.

[124] M. Sekar, M. Pandiaraj, S. Bhansali, N. Ponpandian, and C. Viswanathan, "Carbon fiber based electrochemical sensor for sweat cortisol measurement," *Scientific Reports*, vol. 9, no. 1, p. 403, 2019.

[125] Z. Liu, A. R. Cappola, L. J. Crofford, and W. Guo, "Modeling bivariate longitudinal hormone profiles by hierarchical state-space models," *Journal of the American Statistical Association*, vol. 109, no. 505, pp. 108–118, 2014.

[126] M. A. Lee, N. Bakh, G. Bisker, E. N. Brown, and M. S. Strano, "A pharmacokinetic model of a tissue implantable cortisol sensor," *Advanced Healthcare Materials*, vol. 5, no. 23, pp. 3004–3015, 2016.

[127] H. Raff and T. Carroll, "Cushing's syndrome: From physiological principles to diagnosis and clinical care," *The Journal of Physiology*, vol. 593, no. 3, pp. 493–506, 2015.

[128] R. T. Faghih, "System identification of cortisol secretion: Characterizing pulsatile dynamics," Ph.D. dissertation, Massachusetts Institute of Technology, 2014.

[129] R. T. Faghih, "From physiological signals to pulsatile dynamics: A sparse system identification approach," in *Dynamic Neuroscience*. Springer, 2018, pp. 239–265.

[130] G. Van den Berg, M. Frölich, J. D. Veldhuis, and F. Roelfsema, "Combined amplification of the pulsatile and basal modes of adrenocorticotropin and cortisol secretion in patients with Cushing's disease: Evidence for decreased responsiveness of the adrenal glands," *The Journal Clinical Endocrinology & Metabolism*, vol. 80, no. 12, pp. 3750–3757, 1995.

[131] L. J. Crofford, E. A. Young, N. C. Engleberg, A. Korszun, C. B. Brucksch, L. A. McClure, M. B. Brown, and M. A. Demitrack, "Basal circadian and pulsatile ACTH and cortisol secretion in patients with fibromyalgia and/or chronic fatigue syndrome," *Brain, Behavior and Immunity*, vol. 18, no. 4, pp. 314–325, 2004.

[132] D. D. Pednekar, M. R. Amin, H. F. Azgomi, K. Aschbacher, L. J. Crofford, and R. T. Faghih, "Characterization of cortisol dysregulation in fibromyalgia and chronic fatigue syndromes: A state-space approach," *IEEE Transactions on Biomedical Engineering*, 2020.

[133] D. S. Wickramasuriya and R. T. Faghih, "A cortisol-based energy decoder for investigation of fatigue in hypercortisolism," in *41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, July 2019, pp. 11–14.

[134] V. D'Angelo, G. Beccuti, R. Berardelli, I. Karamouzis, C. Zichi, R. Giordano, M. A. Minetto, M. Maccario, E. Ghigo, and E. Arvat, "Cushing's syndrome is associated with sleep alterations detected by wrist actigraphy," *Pituitary*, vol. 18, no. 6, pp. 893–897, 2015.

[135] L.-U. Wang, T.-Y. Wang, Y.-M. Bai, J.-W. Hsu, K.-L. Huang, T.-P. Su, C.-T. Li, W.-C. Lin, T.-J. Chen, and M.-H. Chen, "Risk of obstructive sleep apnea among patients with Cushing's syndrome: A nationwide longitudinal study," *Sleep Medicine*, vol. 36, pp. 44–47, 2017.

[136] D. M. Arble, G. Copinschi, M. H. Vitaterna, E. Van Cauter, and F. W. Turek, "Chapter 12 - Circadian rhythms in neuroendocrine systems," in *Handbook of Neuroendocrinology*, G. Fink, D. W. Pfaff, and J. E. Levine, Eds. San Diego: Academic

Press, 2012, pp. 271 – 305.

[137] F. Suay and A. Salvador, "Chapter 3 - Cortisol," in *Psychoneuroendocrinology of Sport and Exercise: Foundations, Markers, Trends*, F. Ehrlenspiel and K. Strahler, Eds. Routledge, 2012, pp. 43–60.

[138] Institute of Medicine, *Beyond Myalgic Encephalomyelitis/Chronic Fatigue Syndrome: Redefining an Illness*. Washington, DC: The National Academies Press, 2015.

[139] A. J. Cleare, "The Neuroendocrinology of Chronic Fatigue Syndrome," *Endocrine Reviews*, vol. 24, no. 2, pp. 236–252, 04 2003.

[140] A. S. Papadopoulos and A. J. Cleare, "Hypothalamic–pituitary–adrenal axis dysfunction in chronic fatigue syndrome," *Nature Reviews Endocrinology*, vol. 8, no. 1, p. 22, 2012.

[141] A. Di Giorgio, M. Hudson, W. Jerjes, and A. J. Cleare, "24-hour pituitary and adrenal hormone profiles in chronic fatigue syndrome," *Psychosomatic Medicine*, vol. 67, no. 3, pp. 433–440, 2005.

[142] K. Holtorf, "Diagnosis and treatment of hypothalamic-pituitary-adrenal (HPA) axis dysfunction in patients with chronic fatigue syndrome (CFS) and fibromyalgia (FM)," *Journal of Chronic Fatigue Syndrome*, vol. 14, no. 3, pp. 59–88, 2007.

[143] F. Tanriverdi, Z. Karaca, K. Unluhizarci, and F. Kelestimur, "The hypothalamo–pituitary–adrenal axis in chronic fatigue syndrome and fibromyalgia syndrome," *Stress*, vol. 10, no. 1, pp. 13–25, 2007.

[144] R. G. Krishnan, U. Shalit, and D. Sontag, "Structured inference networks for nonlinear state space models," in *31st AAAI Conference on Artificial Intelligence*, 2017, pp. 2101–2109.

[145] X. Zheng, M. Zaheer, A. Ahmed, Y. Wang, E. P. Xing, and A. J. Smola, "State space LSTM models with particle MCMC inference," *arXiv preprint arXiv:1711.11179*, 2017.