

Hybrid Decoders for Marked Point Process Observations and External Influences

Dilranjan S. Wickramasuriya, *Member, IEEE*, Leslie J. Crofford, Alik S. Widge and Rose T. Faghih, *Member, IEEE*

Abstract—Objective: Internal physiological processes govern multiple state variables within the human body. Estimating these from point process-type bioelectric and biochemical observations is a challenge. Here we seek to estimate cortisol-related energy production and sympathetic arousal based on point process and continuous-valued data while permitting an external influence to affect the state estimates. **Methods:** Traditional point process state-space methods, such as those used for estimating the aforementioned quantities from cortisol and skin conductance measurements respectively, suffer from the inability to permit the state estimates to also fit to an external influence (e.g. labels) or be guided by it. Here we modify an existing recurrent neural network (RNN) approach for state-space estimation through a weighted cost-function to enable a hybrid estimator that has this capability. **Results:** Results on cortisol data based on a hypothetical sleep-wake influence term show how energy production can be estimated by permitting the estimates to fit to the external influence as much as desired. We further show how overfitting may be reduced by using circadian rhythm-based influence terms. Results on skin conductance data also indicate how the method can be used to estimate sympathetic arousal in an experiment containing stressors and relaxation, and permit an external influence as well. **Conclusion:** The RNN-based hybrid method is thus able to recover internal physiological states from point process and continuous-valued observations while permitting an external influence to guide the estimates. **Significance:** The hybrid estimator could be embedded within wearable monitors that can be tailored based on domain expertise or individual feedback.

Index Terms—Biomedical signal processing, state estimation, recurrent neural networks, biomedical monitoring, emotion recognition

The work was supported in part by NSF grants 1942585 – CAREER: MINDWATCH: Multimodal Intelligent Noninvasive brain state Decoder for Wearable Adaptive Closed-loop architectures and 1755780 – CRIL: CPS: Wearable-Machine Interface Architectures. R. T. Faghih served as the senior author.

D. S. Wickramasuriya was with the Electrical and Computer Engineering Department, Cullen College of Engineering, University of Houston. R. T. Faghih is with the Biomedical Engineering Department, Tandon School of Engineering, New York University. She was with the Electrical and Computer Engineering Department, Cullen College of Engineering, University of Houston. L. J. Crofford is with the Division of Rheumatology and Immunology, Vanderbilt University Medical Center. A. S. Widge is with the Department of Psychiatry, University of Minnesota (e-mail: dswickramasuriya@uh.edu, rfaghih@nyu.edu, leslie.j.crofford@vumc.org, awidge@umn.edu). Correspondence should be addressed to senior author R. T. Faghih.

Copyright (c) 2021 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending an email to pubpermissions@ieee.org.

I. INTRODUCTION

PHYIOLOGICAL processes govern multiple internal state variables within the human body. These processes help maintain a stable internal environment despite changes in the external environment and aid the body adapt to the demands imposed thereof (e.g. physical or cognitive loads). Often, the internal state variables are inaccessible and remain difficult to observe directly, but are nonetheless tied to biochemical and bioelectric phenomena that can be measured more easily. These measurements provide a window into estimating the latent physiological state variables of interest.

A. Physiological Point Processes - Cortisol and Skin Conductance

In several instances, the internal states are linked to *point process* observations arising from pulsatile or impulse-like phenomena. The stress hormone cortisol, of which a healthy adult secretes between 15-22 pulses a day, is one such example [1], [2]. Cortisol supplements internal energy production by elevating blood glucose levels as a part of its role in coordinating the stress response [3], [4]. The number, timing and amplitudes of these cortisol pulses reflect changes in the body's internal energy production state. Note that we use the term “energy production” to reflect a utilization of existing glucose reserves rather than an actual production of glucose itself.

Skin conductance, generated by bursts of neural activity to the sweat glands, likewise closely resembles cortisol in its “spikey” appearance and generation dynamics [1], [5]. Skin conductance is a sensitive indicator of sympathetic arousal [6], [7] owing to sweat gland innervation by sympathetic nerve fibers [8]. A fast-varying phasic component, consisting of individual skin conductance responses (SCRs), and a slower-varying tonic component make up a skin conductance signal [9]. SCRs are typically modeled as being generated by bursts of neural activity [5], [10]. SCR rates, amplitudes (to which the *neural impulse* rates and amplitudes are related) and tonic component are three of the most commonly used skin conductance markers of autonomic arousal [11].

A deconvolution algorithm can be used to extract the point process events underlying a cortisol profile or a skin conductance signal (Fig. 1). Due to the cascaded nature of cortisol secretion within the body, the pulses in Fig. 1 can be regarded as abstractions of the CRH (corticotropin-releasing hormone) pulses secreted in the brain [12]. For both cortisol and skin conductance, a marked point process (MPP) accompanied by a

continuous-valued variable constitute the primary observations to which the underlying latent state variable is related (i.e., energy production and sympathetic arousal respectively).




Fig. 1. Deconvolved Blood Cortisol and Skin Conductance Data. A sequence of point process events underlies the generation of a blood cortisol profile and a skin conductance signal. A deconvolution algorithm can be used to extract these underlying pulsatile/impulse events. The upper sub-panel depicts a series of blood cortisol concentrations at 10 min intervals (blue dots), the secretory events (red) and the reconstructed minute-by-minute blood cortisol concentration (black). The lower sub-panel depicts a skin conductance signal (blue), the neural impulses responsible for phasic SCRs (red) and the tonic component (black).

B. State-space Models

Point process state-space methods have been frequently used to estimate latent states tied to observations such those in Fig. 1. These have already found applications in sleep studies [13], neural spiking-based movement and position decoding [14]–[19], behavioral learning [20]–[22], the regulation of anesthesia and comatose states [23]–[25], and the study of heart rate variability [26], [27]. We too have developed several such methods for estimating sympathetic arousal and energy production from skin conductance and cortisol [28]–[34]. The methods utilize Bayesian filtering applied within a larger expectation-maximization (EM) framework.

Traditional such EM-based point process state-space methods have a few notable drawbacks. Firstly, each new observation/feature addition requires a completely new Bayesian filter derivation and EM algorithm (e.g. when extending the model in [21] with one binary and *one* continuous feature to two continuous features to obtain our model in [29]). A second drawback we point out with the aid of an example. If a state-space model for behavioral learning such as described in [20]–[22], [35]–[37] is used to estimate the cognitive learning state of a non-human primate, estimation can only rely on the binary response variables, reaction times etc. since a ground truth for the primate’s learning state is unavailable. If however, sympathetic arousal or energy were to be estimated from a *human* subject based on skin conductance or cortisol,

the subject would be able to provide some feedback (i.e., a form of label) regarding how they felt emotionally or how energetic/lethargic they felt during the experiment (on a rating scale, for instance). While this may not necessarily be the definitive ground truth, it nevertheless does contain some information regarding the unobserved state being estimated.

C. Neural Networks and Hybrid Estimation

We thus raise the question regarding the possibility of estimating a latent state x_k (in this case, arousal or energy) such that x_k also fits to labels. EM-based Bayesian state-space estimation typically does not utilize labels, and it is a challenge to incorporate them into the estimation process. Fitting to labels is traditionally a supervised learning problem. Estimating x_k such that it fits both to physiological observations and labels would be a hybrid form of supervised learning and state estimation. Note that the use of the term “label” in the strict supervised learning sense is somewhat unsuitable for the applications we propose here (we provide an example of x_k fitting to a physiological rhythm later on). We thus choose the broader term “external influence” to denote the additional quantity x_k is permitted to fit to. However, this term too is not synonymous with the external input (usually denoted by u_k) that *drives* x_k in a typical state-space control system.

Recent neural network approaches for state-space estimation (e.g. [38], [39]) present an alternative to traditional EM-based Bayesian state estimation. Methods such as the one in [38] also permit a convenient increase in the number of features without any complete re-derivations. As stated above, it is a challenge for traditional state-space methods to estimate x_k such that it also fits to an external influence (e.g. a label). Here we primarily address this problem by modifying the method in [38] through a hybrid cost function. Our objective is to develop a framework whereby we can estimate energy and sympathetic arousal from MPP type and continuous-valued cortisol and skin conductance features respectively, where the estimated latent state x_k is determined based on both physiological features and an external influence. Our approach can also be adapted to other types of data where the estimated states are required to fit to additional label-type information as well.

II. METHODS

A. Data

1) *Cortisol*: We use the blood cortisol data from the study in [40] for energy state estimation. The study sought to investigate HPA-axis (hypothalamic-pituitary-adrenal) function in patients diagnosed with chronic fatigue syndrome (CFS), fibromyalgia syndrome (FMS) or both. Intravenous blood samples were drawn at 10 min intervals over a 24 h period from patients and matched controls for extracting the blood cortisol levels. Thirty one of the 36 available record pairs were deconvolved by Pednekar *et al.* [41]. Deconvolution yields the pulsatile cortisol profile and other parameters necessary to reconstruct the continuous-valued blood cortisol concentrations at a 1 min time resolution. Thirteen of the 31 patients had CFS (which we label the CFS patient group) while the remaining 18 had either FMS or both FMS and CFS (which we label

the FMS patient group). The first three patients in the FMS group were diagnosed with FMS only while the others had both FMS and CFS.

2) Skin Conductance: We use the skin conductance data from the Non-EEG Biosignals Dataset for Assessment and Visualization of Neurological Status [42] for estimating sympathetic arousal. A number of the records are noise contaminated and we used data from the same subjects that we did in our earlier work in [28]. Similar to [28], we also analyze data from the portion of the experiment containing a period of cognitive stress, relaxation and emotional stress. Each of these lasted for approximately 5 min. A backward-counting task and the Stroop test (a color-word association task) constituted the cognitive stressor. A horror movie clip was used for the emotional stressor. The brief period where instructions for the cognitive tasks were provided is also included and categorized as a stressor [42]. We downsampled the data to 4 Hz prior to lowpass filtering at 0.5 Hz and tonic-phasic separation using cvxEDA [43]. We extracted the neural impulses to the sweat glands from the phasic component using the deconvolution method described in [28].

B. EM-based State-space Estimation

Prior to describing the neural network approach, we briefly review an example EM-based point process state estimation method. We shall subsequently point out certain similarities between the two methods. Assume that our unobserved state variable x_k varies with time following a random walk [28], [34].

$$x_k = x_{k-1} + \varepsilon_k, \quad (1)$$

where $\varepsilon_k \sim \mathcal{N}(0, \sigma_\varepsilon^2)$. Let n_k denote a binary variable indicating the presence or absence of a point process event that x_k gives rise to. Now n_k is Bernoulli-distributed with mass function $p_k^{n_k}(1-p_k)^{1-n_k}$ where $P(n_k=1) = p_k$. Based on the theory of generalized linear models, we can relate x_k to p_k using a logit transformation [44].

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

where β_0 and β_1 are constant coefficients. We have made use of this same formulation in our prior work as well [28]–[34]. Assume we also observe a continuous-valued variable s_k that we take to be related to x_k through

$$s_k = \delta_0 + \delta_1 x_k + w_k, \quad (3)$$

where δ_0, δ_1 are constant coefficients and $w_k \sim \mathcal{N}(0, \sigma_w^2)$. If the regular EM approach were used to estimate x_k and recover the unknown model parameters Θ (consisting of $\delta_0, \delta_1, \sigma_\varepsilon^2$ etc.), we would attempt to maximize [21]

$$Q_1 = \sum_{k=1}^K \mathbb{E} \left[n_k \log(p_k) + (1 - n_k) \log(1 - p_k) \right] - \frac{K}{2} \log(2\pi\sigma_w^2) - \sum_{k=1}^K \frac{\mathbb{E}[(s_k - \delta_0 - \delta_1 x_k)^2]}{2\sigma_w^2}$$

$$- \frac{K}{2} \log(2\pi\sigma_\varepsilon^2) - \sum_{k=1}^K \frac{\mathbb{E}[(x_k - x_{k-1})^2]}{2\sigma_\varepsilon^2} \quad (4)$$

$$= \sum_{k=1}^K \mathbb{E} \left\{ n_k \log \left[\frac{1}{1 + e^{-(\beta_0 + \beta_1 x_k)}} \right] + (1 - n_k) \log \left[1 - \frac{1}{1 + e^{-(\beta_0 + \beta_1 x_k)}} \right] \right\} - \frac{K}{2} \log(2\pi\sigma_w^2) - \sum_{k=1}^K \frac{\mathbb{E}[(s_k - \delta_0 - \delta_1 x_k)^2]}{2\sigma_w^2} - \frac{K}{2} \log(2\pi\sigma_\varepsilon^2) - \sum_{k=1}^K \frac{\mathbb{E}[(x_k - x_{k-1})^2]}{2\sigma_\varepsilon^2}. \quad (5)$$

Here, a Bayesian filter would have to be designed to estimate the expected values for the latent states x_k conditioned on the observations and the parameter estimates. While this remains the standard approach, and we too have developed several such estimators [28]–[34], it nonetheless suffers from the limitations described above. Namely, an inflexibility with regard to feature addition/modification (e.g. each addition necessitates a completely new Bayesian filter derivation and EM algorithm) and a difficulty in incorporating a mechanism whereby x_k can be permitted to have an external influence affect its estimation. In contrast, adding new features to [38] requires no complete re-derivation and also affords a means whereby x_k can be estimated by permitting an external label-like influence.

C. Neural Networks and Training

Let us now consider the neural network approach in [38]. Here, the general Gaussian state-space model

$$x_k \sim \mathcal{N}(f_{\mu_x}(x_{k-1}), f_{\sigma_x^2}(x_{k-1})) \quad (6)$$

$$y_k \sim \Pi(f_y(x_k)), \quad (7)$$

where y_k represents the observations, is assumed. 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 .

We formally describe hybrid estimation based on a modification of the method in [38] in subsection III-A below. In [38], the neural networks are trained by maximizing a probability term. Here we use a re-weighted cost function that is a combination of the negative of this term and an additional penalization term based on the external influence which we denote by l_k . We use ρ to denote the coefficient determining how much weight is assigned to either of these terms [this will be formally defined in (12)]. We evaluate our proposed hybrid estimation scheme as described below.

(i) Cortisol – We do not possess additional information from the subjects regarding how energetic/lethargic they felt (i.e., similar to a label) during the course of the 24 h periods. Therefore, we generated a hypothetical l_k to illustrate the possibility of hybrid estimation. Since cortisol-related energy production is somewhat lower towards nighttime than it is during the morning, we generated a binary-like l_k term to

match the sleep-wake times of the subjects enrolled in the study [41]. We first trained the neural networks for 500 epochs with $\rho = 0$ (i.e., disallowing the external influence). We then took this pre-trained model, set $\rho = 0.5$, introduced l_k and trained for several more epochs.

While hybrid estimation with the use of a binary-like sleep-wake l_k term largely serves an illustrative purpose in this case, we then present a more practical application of using l_k to reduce overfitting to the continuous-valued variable. Now cortisol is known to exhibit circadian variation [45]. Therefore, we generated separate circadian l_k terms for each subject and used them to undo the overfitting. To calculate these subject-specific l_k 's, we first took the 500-epoch pre-trained models ($\rho = 0$) for each one, and then calculated a two-harmonic circadian based on a least-squares fit to each subject's x_k estimate. The two-harmonic l_k term is of the form [46]

$$l_k = a_0 + \sum_{i=1}^2 \left[a_i \sin\left(\frac{2\pi i k}{1440}\right) + b_i \cos\left(\frac{2\pi i k}{1440}\right) \right], \quad (8)$$

and therefore only the a_i and b_i coefficients need to be calculated for each subject using least squares. We then set $\rho = 0.75$ and trained for an additional $j = 20$ epochs.

(ii) Skin Conductance – We estimate sympathetic arousal from skin conductance using data from the experiment in [42] containing different stress periods interspersed by relaxation. Similar to [28], we calculate a high arousal index (HAI) as the probability that $x_k > x_{\text{thresh}}$ as well. The threshold x_{thresh} is selected as the median state value [28]. The HAI is analogous to the ideal observer certainty level in [20], [21] and expresses the probability of x_k exceeding a certain baseline (since p_k is related to x_k , the term can also be calculated based on p_k exceeding an equivalent baseline). We calculate the state estimates for all the subjects with $\rho = 0$. As an illustration of the possibility of hybrid estimation with skin conductance data, we also select a hypothetical square wave-like l_k term and obtain the resulting estimate when x_k is made to fully conform to l_k .

For cortisol data, we use a neural network configuration with two layers in the SSNN and a hidden layer dimensionality of 256. We also select an RNN size of 256. For skin conductance, we selected an SSNN hidden layer dimensionality of 200 with two layers and an RNN size of 200. Effects of changing the RNN size, the number of layers in the SSNN and the SSNN hidden layer dimensionality on the state estimates are described in more detail in the supplementary information.

III. RESULTS

A. Model

Taking ψ and ϕ to denote the state-space model parameters and the RNN parameters respectively, the networks are trained by maximizing

$$\begin{aligned} & \sum_{k=1}^K \mathbb{E}_{q_\phi(x_k|\vec{y})} [\log p_\psi(y_k|x_k)] - \text{KL}(q_\phi(x_1|\vec{y})||p_\psi(x_1)) \\ & - \sum_{k=2}^K \mathbb{E}_{q_\phi(x_{k-1}|\vec{y})} [\text{KL}(q_\phi(x_k|x_{k-1},\vec{y})||p_\psi(x_k|x_{k-1}))], \end{aligned} \quad (9)$$

where $p_\psi(\cdot)$ and $q_\phi(\cdot)$ denote density functions [38]. The actual training is performed within the algorithm as a minimization of the negative term. We label this *negative* term Q_2 . Analogous to the traditional EM method, in this new approach: (i) the SSNN replaces the explicit state-space model [such as described in (1)-(3)]; (ii) the RNN replaces the Bayesian filter for state estimation; (iii) the weights of the neural networks replace Θ . Since neural networks are used to learn the state-space model, more complicated state transitions and input-output relationships are permitted. These, however, come at the expense of interpretability. The neural network weights parameterized by ψ and ϕ are updated using stochastic backpropagation and details of this can be found in [38].

We now point out some general similarities with the traditional EM method, particularly with regard to the terms in Q_1 and Q_2 depending on the type of variables that are present. For instance, if a binary variable n_k is present among the observations y_k , Q_2 contains

$$-\sum \left[n_k \log \left(\frac{1}{1 + e^{-f_n(x_k)}} \right) + (1 - n_k) \log \left(1 - \frac{1}{1 + e^{-f_n(x_k)}} \right) \right], \quad (10)$$

similar to the first term in (5), except that now $f_n(\cdot)$ is a function learned by the SSNN. Similarly, if a continuous-valued variable s_k is present in y_k , there is the term

$$\sum \frac{1}{2} \log [2\pi f_{\sigma_s^2}(x_k)] + \frac{[s_k - f_{\mu_s}(x_k)]^2}{2f_{\sigma_s^2}(x_k)} \quad (11)$$

in Q_2 , where $f_{\mu_s}(\cdot)$ and $f_{\sigma_s^2}(\cdot)$ represent mean and variance functions learned by the SSNN. Again, note the similar term involving s_k in (5).

B. Hybrid Estimation

In our case, the observations consist of an MPP and a continuous variable. We take n_k to denote the presence or absence of a point process event (cortisol pulse or neural impulse), r_k to denote the event amplitude (mark) and s_k to denote the continuous variable (blood cortisol concentration or tonic level). We make the following two modifications to the original method in [38] to enable hybrid estimation:

(i) We first change the cost function to

$$\tilde{Q}_2 = (1 - \rho)Q_2 + \rho \sum (x_k - l_k)^2, \quad (12)$$

where $0 \leq \rho \leq 1$. Note carefully the form of (12). Typical supervised learning problems (e.g. regression) have a cost function similar to the second term on the right of (12), where l_k is a form of label or known ground truth. On the other hand, traditional EM-based Bayesian state-space estimation usually only has a term similar to Q_2 . By changing the cost function to the weighted sum in (12), we enable a hybrid of the two approaches. Our general method is to first train the neural networks with $\rho = 0$ similar to the original method in [38], then take the pre-trained model, choose a positive value for ρ , introduce l_k and re-train for several more epochs to modify the state estimates. This enables a class of hybrid estimators




Fig. 2. Illustration of the Effect of the External Influence Term l_k in Cortisol-based Energy State Estimation. The sub-figures respectively depict the energy state estimates for CFS patient 1 from a 500-epoch pre-trained model with $\rho = 0$ (left), after training an for additional $j = 25$ epochs with $\rho = 0.5$ (center) and after training an for additional $j = 100$ training epochs with $\rho = 0.5$ (right). 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 sleep-wake l_k term (blue).

that allows x_k to be estimated based on a mix of physiological data and an external influence.

(ii) Secondly, similar to [34], we change the form of (11) to sum only over the locations where the point process events occur for the MPP amplitudes r_k .

C. Cortisol

Deconvolving the data in [41] yields the pulsatile cortisol profile (n_k, r_k) and the cortisol infusion and clearance rates necessary to reconstruct the cortisol concentration s_k at a 1 min resolution [41]. The neural networks yield the state estimates $x_k \sim \mathcal{N}(x_{k,NN}, \sigma_{k,NN}^2)$ and we provide $x_{k,NN}$ in the figures with confidence intervals calculated using $\sigma_{k,NN}^2$.

Fig. 2 shows the effect of introducing a hypothetical binary-like sleep-wake l_k term and the entailing (gradual) conformity of x_k to l_k during state estimation. With $\rho = 0$, x_k is only estimated from the data. However, after setting $\rho = 0.5$ and training for $j = 25$ more epochs, the shape of x_k gradually begins to conform to l_k . After $j = 100$ epochs, x_k completely fits to l_k . By changing ρ and the number of additional training epochs j , the external influence term l_k can be permitted to affect the state estimates as much as desired.

The practical application of the use of l_k in reducing overfitting to s_k by means of a circadian term is shown for a single subject (CFS patient 1) in Fig. 3 (this overfitting to s_k can especially be seen in the state estimates for the pre-trained model with $\rho = 0$ in Fig. 2). The results for all 31 subject pairs using this overfitting control technique are provided in the supplementary information (note that the neural networks are trained for each subject separately). Since the choice of $\rho = 0.75$ and $j = 20$ to undo overfitting may seem arbitrary, additional results based on varying these parameters and a further discussion of them is provided in

the supplementary information. In general, the number of additional training epochs j , the coefficient ρ and the l_k term can all be customized based on the degree of overfitting control that is desired. Taking the pre-trained model, and then training for several additional epochs with l_k included and $\rho > 0$ forces x_k to be influenced externally undoing some of the overfitting to s_k . We do, however, note one drawback with this method of overfitting control. When training with the l_k term included and $\rho > 0$, we are forcing the neural networks to fit x_k to not just the physiological data alone, but rather to an external influence term as well. Consequently, as x_k deviates away from its original fit and gradually begins to conform to l_k , the fits to the *other* physiological features (i.e., to the MPP amplitudes r_k and the continuous-valued feature s_k) tend to become noisy. This effect can be seen in Figs. 2 and 3.

There is a scarcity of point process-related state estimation work in the literature for pulsatile hormones such as cortisol. It is of note that even a system-theoretic method for deconvolving serum cortisol data to estimate the underlying pulsatile secretion timings and amplitudes with person-specific model parameters was only developed comparatively recently [1], [12]. Our own estimation work in [33] also only utilized simulated data. Now label-like information is unavailable for the experimental cortisol data considered here [40]. Consequently, and due to the relatively high number of subjects in our dataset (31 subject pairs), the evaluation of the energy state estimation results is largely qualitative in nature.

We note two primary features of interest in the qualitative evaluation of our results. Firstly, energy state estimates are consistent with the cortisol awakening response (CAR) [47] across subjects. As stated earlier, cortisol secretion is circadian and is synchronized to light-dark and sleep-wake cycles [40]. Typically, cortisol begins to rise towards early morning and




Fig. 3. Cortisol-based Energy State Estimation with a Circadian l_k Term Included to Reduce Overfitting. 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).

reaches peak values shortly after awakening. This is known as the CAR. Taking the time interval between 11:00 p.m. and 8:00 a.m. as the approximate sleep-wake period (similar to [41]), a corresponding rise in the cortisol-related energy states towards morning awakening is clearly visible across subjects. This general trend occurs for both patients and controls with lesser inter-subject variations. A second qualitative feature is the diminishing of the energy state towards evening, with the lowest values typically reached during sleep late at night. This too is in accordance with the known physiology of cortisol [48], [49] and can be seen for both patients and controls.

We next evaluate group-level differences. Now CFS is a chronic condition of unknown medical etiology where patients are primarily characterized by severe fatigue [50]. Pain, post-exertional malaise and difficulties in sleeping and concentrating are also symptomatic of CFS patients [50]. FMS is a similar condition where patients experience comparable symptoms [51]. Findings have been mixed with regard to cortisol differences between CFS and FMS patients and matched controls [52]–[55]. A few serial sample studies have indicated a lack of significant mean differences between patients and controls [40], [53] (here we use the data from [40]), and it currently remains unclear whether in fact definite cortisol differences do exist in these conditions, and if any such changes are causal or consequential. We provide a comparison of the averaged x_k and p_k values for all patients and matched controls in the supplementary information. However, caution

must be exercised when interpreting these averaged values for two reasons. Firstly, the neural networks can learn p_k estimates that are inverted for some subjects. These then cancel out with non-inverted p_k estimates from other subjects. Secondly, the x_k estimates can also lie in slightly different ranges for different subjects. However, we too observe that there do not appear to be any significant differences in the x_k or p_k estimates between patients and controls.

D. Skin Conductance

Skin conductance-based sympathetic arousal estimates for the subjects considered from the dataset in [42] with $\rho = 0$ are provided in the supplementary information. We also perform hybrid estimation with a hypothetical square wave-like l_k term for participant 1 to serve an illustrative purpose. This result, along with the corresponding state estimate with $\rho = 0$ is shown in Fig. 4.

The data from [42] considered here contains two psychological stressors, the first of which is cognitive and the second emotional. The cognitive stressor, which again consisted of two parts, required the participants to count backwards in sevens beginning at 2485 and then perform the Stroop test. The emotional stressor consisted of watching a clip from a zombie movie [42]. In the case of $\rho = 0$, the sympathetic arousal estimates generally agree with those from our previous state-space models in [28], [29], [31]. With the exception of participant 4, arousal remains high during cognitive stress for all participants. The active engagement required to perform the cognitive tasks coupled together with the alerting of mistakes through a buzzer [42] likely elicited a sympathetic stress response in them. Arousal gradually declines thereafter as the relaxation period commences. The rate of decline however, shows inter-subject variability due to individual differences in the stress response and recovery mechanism. The horror movie clip fails to elicit as much stress from the participants although a brief increase is seen in x_k at the start. The tonic s_k levels also tend to be lower during the clip. Watching the clip, unlike the earlier stressor, only involves a passive type of engagement and likely did not generate a sufficient fear response as anticipated. Participant 4 is an exception for whom arousal is significantly higher during relaxation. A sharp increase in skin conductance is seen for this participant as the relaxation period begins. This is unusual, and may possibly have been due to a motion artifact causing the increase rather than the actual onset of relaxation itself.

While estimates from the RNN method generally agree with results from earlier state-space work in cortisol and skin conductance, three drawbacks are to be noted. Firstly, small fluctuations remain in the x_k estimates compared to those in [28], [29], [31], [33]. The effect of these fluctuations becomes even more pronounced in the HAI values when x_k is used for a further probability calculation. The neural networks have thousands of weights to estimate while the datasets are small. This may likely be why some of the estimates are noisy. Physiological data collection can be a resource-expensive process (e.g. blood sample acquisition at 10 min intervals over a 24 h period), and this highlights one of the




Fig. 4. Skin conductance-based Sympathetic Arousal Estimation. The sub-figures depict arousal estimation without (left) and with (right) a hypothetical l_k term included. 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 sympathetic arousal state x_k (purple) with its 95% confidence limits; the hypothetical l_k term (blue) is shown in the right sub-figure; (f) the 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). The counting task and the Stroop test formed the cognitive stress portion of the experiment.

drawbacks of using neural networks to perform estimation with limited data. Secondly, the x_k estimates do appear to overfit to the continuous-valued variable s_k in the case of both cortisol and skin conductance. With cortisol, we were able to undo overfitting by means of a circadian l_k term. For skin conductance, however, individual stress responses can vary significantly and it remains unclear what type of l_k should be used to undo overfitting but retain individual variability. The choice of l_k thus remains a topic for further research here. Thirdly and finally, the estimates of p_k can be inverted for certain subjects with the RNN method (the neural networks appear to be learning an inverse of the point process probability here). The same inversion can also occur for x_k (discussed in the supplementary information). Unfortunately, it is a challenge to guide the neural networks to consistently learn p_k 's and x_k 's that are non-inverted. Nevertheless, the hybrid RNN-based approach does permit x_k to be influenced externally – a challenge for conventional state-space methods.

IV. DISCUSSION

Many physiological states within the human body remain unobserved. Nevertheless, the biochemical and bioelectric phenomena they are tied to provide a means for their estimation. Point process state-space methods, such as the ones we have developed in our previous work [28]–[34], can be used for this purpose. Unfortunately, these methods do not conveniently permit an external influence such as labels or medical domain

expertise to affect or guide the states being estimated. Here we have illustrated how an existing neural network approach for state estimation can be modified to address this shortcoming. We provide a discussion related to certain aspects common to both cortisol and skin conductance-based estimation below.

A. Data Acquisition

The cortisol data used in our work was acquired from subjects under constrained settings. Ideally, however, data could be collected in an unconstrained or free-range setting, and would also enable more samples to be acquired. Since meal times and sleep affect the secretion of circadian hormones such as cortisol, it would be necessary for subjects to maintain records of their daily activities in such an experiment. The collection of serum cortisol, instead of salivary cortisol, would also prove an additional challenge here. Feedback from subjects regarding symptomatic events or fatigue levels (i.e., labels) could be acquired through wearable devices. Salivary cortisol, while affording convenience, does not provide as clear a picture of basal secretion compared to its serum counterpart. Future work, however, would involve developing a deconvolution algorithm and energy state estimator based on salivary hormonal samples due to the comparative ease of data collection. Similar real-world skin conductance data acquisition would also be advantageous for testing out our methods in unconstrained settings.

B. Validation Challenges and Generalization

The lack of a ground truth during state estimation presents a number of difficulties, particularly with regard to the validation of the final estimates and the choice of features. This, however, is an issue common to other latent physiological state estimation problems as well. A discussion on issues related to feature selection, validation and the selection of the level of overfitting control is thus provided in the supplementary information.

Traditionally, the data is split into training, validation and test sets when using machine learning. However, we do not use machine learning in the traditional sense here. Inter-subject variability is a known issue in biomedical research. Our use of neural networks for state-space estimation is more along the lines of determining the best set of weights that need to be programmed into a particular patient's personalized healthcare device. Given individual variability and the need for personalization, a patient's device would require customization based on his/her *own* data. However, since this can raise issues with performance on unseen data, we have provided some additional results with held-out data here and in the supplementary information. For instance, Fig. 5 shows two cases where x_k for CFS patient 1 is estimated using: (i) the model trained with data from the other patients in the CFS group (i.e., holding out CFS patient 1's data); (ii) the circadian-based overfitting controlled model trained for FMS patient 1. The estimates do appear reasonable and help support the model's viable performance on hitherto unseen data.

Physiological changes occur in the human body over time. Such changes can also occur due to disease onset, or in the case of patients, with changes in medication dosages, types or conditions. The neural networks here learn how a particular latent physiological state x_k evolves with time, what its relationship is to the observed MPP and continuous-valued features and how it is to be estimated. The weights of the neurons thus capture physiological information. If changes occur in an individual, with no corresponding update in the neuron weights, it is likely that the accuracy of the x_k estimates and the performance of the system will diminish. We thus propose that the neural networks be re-trained periodically (for instance, if a change in the medical condition of an individual were to occur). Having the neural networks be person-specific (as we have done here) as well as conducting periodic re-training would help the method generalize both to the individual (e.g. as may be applicable in the case of medical conditions/genetics unique to a person) and generalize over time but to the same individual.

C. The Choice of External Influence and Degree of Fit

As evident from (12) and the results discussed above, the latent state x_k can not only be estimated based on physiological observations and an external influence l_k , it can also be made to completely ignore either one of them. Consequently, and in the absence of a ground truth, it is a valid concern as to how trustworthy estimates of x_k can be obtained in different application contexts. Therefore, l_k must remain physiologically plausible in order for the final estimates to be reliable. This

then, is bound up with two further questions: (i) What should l_k itself be in different contexts? (ii) How much should x_k be permitted to fit to it in each case? Our formulation does not impose a restriction on the nature of l_k and the question of how it is to be selected will likely require further research. In the meantime, we propose that l_k be chosen in one of two ways to ensure reliable state estimates. Firstly, l_k can be chosen based on medical expertise, domain knowledge well-established in the literature (e.g. biological rhythms, rule sets) or based on models developed with insight from these sources. Secondly, l_k can be chosen based on individual feedback (which would allow for personalization). An arbitrary selection of l_k will likely not provide any guarantee on the reliability of the x_k estimates. Note that the choice of l_k also affects interpretability. If l_k were chosen based on either of the two options just stated, we could interpret the swaying away of x_k towards l_k as permitting domain knowledge or subject-specific personalization to affect the estimates. Interpretability however, would be lost if x_k were to be influenced by an arbitrary l_k .

Moreover, in each such case, the degree to which x_k should be permitted to fit to l_k will likely require further research and experimentation as well. This degree of fit is related to the choice of ρ and the number of additional training epochs j , and will also vary based on what l_k is. A possible future direction in this regard is to obtain data where skin conductance *and* subject-provided labels are available (e.g. [56], [57]) and determine, in general, to what extent x_k needs to be permitted to fit to l_k . Similar experiments could also be performed for cortisol based on subject-provided feedback of how energetic/lethargic they felt. Thus both the choice l_k and selecting its degree of influence on x_k will likely require further research.

D. Application Scenarios

The type of hybrid estimator we present may find applications in a number of scenarios. It could, for instance, be used in wearable monitors. The estimator could also be embedded within a larger control loop for regulating energy or emotional arousal in patients with certain cortisol or neuropsychiatric disorders. In each case the choice of ρ permits a certain degree of flexibility in how much l_k is allowed to influence x_k – a choice that would likely have to be taken by a physician or the wearer. Consider, for instance, a patient with post-traumatic stress disorder (PTSD) or depression (known to involve symptoms of either hyperarousal or abnormally low arousal levels respectively [28]), fitted with a smart skin conductance-based wearable device, being monitored for a period of time. The patient may periodically have to fill out questionnaires with details of how he/she is feeling on an emotional arousal rating scale. In the meantime, the device's estimates would also provide continuous arousal levels. The value of ρ would have to be tuned to obtain a complete picture of the patient's emotional state based on the two hybrid sources of information. Note also that the implementation architecture and frequency of model re-training would depend on the application scenario and not have a "one-size-fits-all"




Fig. 5. Cortisol-based Energy State Estimation with Held-out Data. The sub-figures depict the energy estimates for CFS patient 1 with the neural network model trained using data from the other patients in the CFS group (left) and with the circadian-based overfitting controlled neural network model trained FMS patient 1 (right). 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.

approach. Domain knowledge regarding circadian rhythms do not change rapidly and hence the models may not require frequent re-training. The scenario may be different for arousal perhaps with more rapidly-changing emotional environments.

This type of hybrid or labeled estimator may also have utility in wearables or other remote monitoring devices for clinical scenarios, particularly those where prior information or expectations are available. In one example, PTSD, anxiety, and related mental disorders can lead to dysregulation of stress responses, which may lead to abnormal skin conductance dynamics. These may only be apparent, however, when patients are exposed to stressful situations. It may be possible to provide information on the probability of such stressors through our labeling technique. For instance, there are usually fewer stressors or trauma reminders in a patient's home than in work, school, or other environments far from home. In a wearable monitor, the patient's distance from their home could be tracked by GPS, and could be provided as an input/label. This would effectively act to shift the prior probability of the stress state to a different level. Similarly, in depression, patients can have dysregulated cortisol secretion, e.g. a loss of normal patterns of circadian variation [58]. If the expected circadian rhythm can be provided as input to a monitoring system, then deviations from that rhythm will be more readily detectable. In that application, the error terms/gain terms in the estimator can be as important as the tracked state itself. Both these models could then be used to track response to treatment or relapse of illness, providing guidance to clinicians.

E. Alternate Methods and Future Work

Here we have used neural networks for estimation with the hybrid formula in (12). However, we could alternatively have used EM-based Bayesian filters as well by introducing a weighted l_k -based cost term in addition to the conventional log-likelihood term. Then, instead of solving closed-form expressions for all the model parameters Θ as we usually do in [28]–[34], we could solve for Θ at once through an optimization formulation. This would permit a class of hybrid Bayesian filters where labels could influence the estimates. Nevertheless, jointly determining the model parameters through such an optimization will likely be challenging, perhaps even more so in the case of models such as [29], [30] where a larger number of parameters are present.

Certain types of automated hormone infusion pumps (e.g. insulin pumps) are currently available in the commercial market as are devices for monitoring biochemical concentrations within the body. While these pumps largely function based on monitoring a single variable (e.g. blood glucose), it is likely that chemical concentrations within the body are, in reality, affected by more than a single quantity. Cognitive stress, for instance, affects both energy production and glucose levels. While further research would be necessary, it is possible that next-generation infusion pumps could monitor multiple variables such as skin conductance (related to stress) and cortisol simultaneously for adjusting dosages automatically. In our work here we have estimated separate psychological and hormone-related energy states. However, the observations

could be combined to simultaneously obtain a more comprehensive view of the human body and be embedded within a closed-loop control framework (e.g. for energy regulation based on hormones and psychological stress levels). The concept of energy may also be multi-faceted. For instance, while a person involved in manual labor may experience physical exhaustion, another individual may experience cognitive fatigue due to stressful work demands or the skills/concentration required to perform a task. Certain other types of measurements (e.g. heart rate variability) could be explored to determine factors that correlate with different types of fatigue to obtain a holistic view of energy production within the human body.

V. CONCLUSION

Changes in internal states within the human body give rise to electrical and chemical phenomena that can be easily measured using sensors. While it is possible to estimate these latent states purely from sensor data, it is likely that feedback provided by subjects, medical diagnostic data etc. also contain information regarding the states being estimated, and should thus be permitted to influence/guide state estimation. Consequently, latent physiological state estimation may need to be based on a hybrid fusion of sensor data and external information. Here we have illustrated how latent states tied to markers in skin conductance and cortisol can be estimated based on physiological data and as well as an external influence. This work functions as a proof-of-principle of a hybrid form of supervised learning and state estimation. Our approach could also be applied to other types of point process data and have important applications to personalized wearable monitoring.

REFERENCES

- [1] R. T. Faghih *et al.*, "Deconvolution of serum cortisol levels by using compressed sensing," *PLoS One*, vol. 9, no. 1, p. e85204, 2014.
- [2] —, "An optimization formulation for characterization of pulsatile cortisol secretion," *Frontiers in Neuroscience*, vol. 9, p. 228, 2015.
- [3] A. Clow *et al.*, "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>
- [4] 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.
- [5] 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.
- [6] H. D. Critchley, "Electrodermal responses: What happens in the brain," *The Neuroscientist*, vol. 8, no. 2, pp. 132–142, 2002.
- [7] M.-Z. Poh *et al.*, "A wearable sensor for unobtrusive, long-term assessment of electrodermal activity," *IEEE Transactions on Biomedical Engineering*, vol. 57, no. 5, pp. 1243–1252, 2010.
- [8] P. A. Low, "Chapter 51 - Sweating," in *Primer on the Autonomic Nervous System (Third Edition)*, 3rd ed., D. Robertson, I. Biaggioni, G. Burnstock, P. A. Low, and J. F. Paton, Eds. San Diego: Academic Press, 2012, pp. 249 – 251.
- [9] M. Benedek and C. Kaernbach, "Decomposition of skin conductance data by means of nonnegative deconvolution," *Psychophysiology*, vol. 47, no. 4, pp. 647–658, 2010.
- [10] M. R. Amin and R. T. Faghih, "Robust inference of autonomic nervous system activation using skin conductance measurements: A multi-channel sparse system identification approach," *IEEE Access*, vol. 7, pp. 173 419–173 437, 2019.
- [11] S. D. Kreibitz, "Autonomic nervous system activity in emotion: A review," *Biological Psychology*, vol. 84, no. 3, pp. 394–421, 2010.
- [12] R. T. Faghih, "System identification of cortisol secretion: Characterizing pulsatile dynamics," Ph.D. dissertation, Massachusetts Institute of Technology, 2014.
- [13] M. J. Prerau *et al.*, "Tracking the sleep onset process: An empirical model of behavioral and physiological dynamics," *PLoS Computational Biology*, vol. 10, no. 10, p. e1003866, 2014.
- [14] M. M. Shanchei *et al.*, "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.
- [15] —, "Neural population partitioning and a concurrent brain-machine interface for sequential motor function," *Nature Neuroscience*, vol. 15, no. 12, p. 1715, 2012.
- [16] —, "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.
- [17] X. Deng *et al.*, "Clusterless decoding of position from multiunit activity using a marked point process filter," *Neural Computation*, vol. 27, no. 7, pp. 1438–1460, 2015.
- [18] K. Arai *et al.*, "Marked point process filter for clusterless and adaptive encoding-decoding of multiunit activity," *bioRxiv*, p. 438440, 2018.
- [19] A. Yousefi *et al.*, "Real-time point process filter for multidimensional decoding problems using mixture models," *bioRxiv*, p. 505289, 2018.
- [20] A. C. Smith *et al.*, "Dynamic analysis of learning in behavioral experiments," *Journal of Neuroscience*, vol. 24, no. 2, pp. 447–461, 2004.
- [21] M. J. Prerau *et al.*, "Characterizing learning by simultaneous analysis of continuous and binary measures of performance," *Journal of Neurophysiology*, vol. 102, no. 5, pp. 3060–3072, 2009.
- [22] T. P. Coleman *et al.*, "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.
- [23] Y. Yang and M. M. Shanchei, "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.
- [24] Y. Yang *et al.*, "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.
- [25] Y. Yang and M. M. Shanchei, "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.
- [26] 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.
- [27] —, "Analysis of heartbeat dynamics by point process adaptive filtering," *IEEE Transactions on Biomedical Engineering*, vol. 53, no. 1, pp. 4–12, 2006.
- [28] D. S. Wickramasuriya *et al.*, "Skin conductance as a viable alternative for closing the deep brain stimulation loop in neuropsychiatric disorders," *Frontiers in Neuroscience*, vol. 13, p. 780, 2019.
- [29] 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.
- [30] —, "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.
- [31] D. S. Wickramasuriya *et al.*, "A state-space approach for detecting stress from electrodermal activity," in *40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, July 2018, pp. 3562–3567.
- [32] 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.
- [33] —, "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.
- [34] 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.
- [35] A. Yousefi *et al.*, "COMPASS: An open-source, general-purpose software toolkit for computational psychiatry," *Frontiers in Neuroscience*, vol. 12, p. 957, 2019.

- [36] —, “Cognitive state prediction using an EM algorithm applied to Gamma distributed data,” in *37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE, 2015, pp. 7819–7824.
- [37] —, “Predicting learning dynamics in multiple-choice decision-making tasks using a variational Bayes technique,” in *39th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE, 2017, pp. 3194–3197.
- [38] R. G. Krishnan *et al.*, “Structured inference networks for nonlinear state space models,” in *31st AAAI Conference on Artificial Intelligence*, 2017, pp. 2101–2109.
- [39] X. Zheng *et al.*, “State space LSTM models with particle MCMC inference,” *arXiv preprint arXiv:1711.11179*, 2017.
- [40] L. J. Crofford *et al.*, “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.
- [41] D. D. Pednekar *et al.*, “Characterization of cortisol dysregulation in fibromyalgia and chronic fatigue syndromes: A state-space approach,” *IEEE Transactions on Biomedical Engineering*, 2020.
- [42] J. Birjandtalab *et al.*, “A non-EEG biosignals dataset for assessment and visualization of neurological status,” in *IEEE International Workshop on Signal Processing Systems (SiPS)*, 2016, pp. 110–114.
- [43] A. Greco *et al.*, “cvxEDA: A convex optimization approach to electrodermal activity processing,” *IEEE Transactions on Biomedical Engineering*, vol. 63, no. 4, pp. 797–804, 2016.
- [44] P. McCullagh and J. A. Nelder, *Generalized Linear Models*. CRC press, 1989, vol. 37.
- [45] E. D. Weitzman *et al.*, “Twenty-four hour pattern of the episodic secretion of cortisol in normal subjects,” *The Journal of Clinical Endocrinology & Metabolism*, vol. 33, no. 1, pp. 14–22, 1971.
- [46] E. N. Brown *et al.*, “A stochastic differential equation model of diurnal cortisol patterns,” *American Journal of Physiology-Endocrinology and Metabolism*, vol. 280, no. 3, pp. E450–E461, 2001.
- [47] A. Clow *et al.*, “The awakening cortisol response: Methodological issues and significance,” *Stress*, vol. 7, no. 1, pp. 29–37, 2004.
- [48] D. M. Arble *et al.*, “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. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/B9780123750976100125>
- [49] 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.
- [50] Institute of Medicine, *Beyond Myalgic Encephalomyelitis/Chronic Fatigue Syndrome: Redefining an Illness*. Washington, DC: The National Academies Press, 2015.
- [51] D. Buskila and J. Press, “Neuroendocrine mechanisms in fibromyalgia–chronic fatigue,” *Best Practice & Research Clinical Rheumatology*, vol. 15, no. 5, pp. 747 – 758, 2001. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1521694201901913>
- [52] A. J. Cleare, “The neuroendocrinology of chronic fatigue syndrome,” *Endocrine Reviews*, vol. 24, no. 2, pp. 236–252, 04 2003. [Online]. Available: <https://doi.org/10.1210/er.2002-0014>
- [53] 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.
- [54] 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.
- [55] F. Tanriverdi *et al.*, “The hypothalamo–pituitary–adrenal axis in chronic fatigue syndrome and fibromyalgia syndrome,” *Stress*, vol. 10, no. 1, pp. 13–25, 2007.
- [56] S. Koelstra *et al.*, “DEAP: A database for emotion analysis using physiological signals,” *IEEE Transactions on Affective Computing*, vol. 3, no. 1, pp. 18–31, 2012.
- [57] M. Soleymani *et al.*, “A multimodal database for affect recognition and implicit tagging,” *IEEE Transactions on Affective Computing*, vol. 3, no. 1, pp. 42–55, 2012.
- [58] G. W. Arana *et al.*, “The dexamethasone suppression test for diagnosis and prognosis in psychiatry: commentary and review,” *Archives of General Psychiatry*, vol. 42, no. 12, pp. 1193–1204, 1985.