



Published in final edited form as:

Inverse Probl. 2022 May ; 38(5): . doi:10.1088/1361-6420/ac5ac7.

Blood and Breath Alcohol Concentration from Transdermal Alcohol Biosensor Data: Estimation and Uncertainty Quantification via Forward and Inverse Filtering for a Covariate-Dependent, Physics-Informed, Hidden Markov Model

Clemens Oszkinat¹, Tianlan Shao¹, Chunming Wang¹, I. G. Rosen^{1,§}, Allison D. Rosen², Emily B. Saldich³, Susan E. Luczak³

¹Department of Mathematics, University of Southern California, Los Angeles, CA 90089, USA

²Department of Epidemiology, Fielding School of Public Health, University of California, Los Angeles, CA 90095, USA

³Department of Psychology, University of Southern California, Los Angeles, CA 90089, USA

Abstract

Transdermal alcohol biosensors that do not require active participation of the subject and yield near continuous measurements have the potential to significantly enhance the data collection abilities of alcohol researchers and clinicians who currently rely exclusively on breathalyzers and drinking diaries. Making these devices accessible and practical requires that transdermal alcohol concentration (TAC) be accurately and consistently transformable into the well-accepted measures of intoxication, blood/breath alcohol concentration (BAC/BrAC). A novel approach to estimating BrAC from TAC based on covariate-dependent physics-informed hidden Markov models with two emissions is developed. The hidden Markov chain serves as a forward full-body alcohol model with BrAC and TAC, the two emissions, assumed to be described by a bivariate normal which depends on the hidden Markovian states and person-level and session-level covariates via built-in regression models. An innovative extension of hidden Markov modeling is developed wherein the hidden Markov model framework is regularized by a first-principles PDE model to yield a hybrid that combines prior knowledge of the physics of transdermal ethanol transport with data-based learning. Training, or inverse filtering, is effected via the Baum-Welch algorithm and 256 sets of BrAC and TAC signals and covariate measurements collected in the laboratory. Forward filtering of TAC to obtain estimated BrAC is achieved via a new physics-informed regularized Viterbi algorithm which determines the most likely path through the hidden Markov chain using TAC alone. The Markovian states are decoded and used to yield estimates of BrAC and to quantify the uncertainty in the estimates. Numerical studies are presented and discussed. Overall good agreement between BrAC data and estimates was observed with a median relative peak error of 22% and a median relative area under the curve error of 25% on the test set. We also demonstrate that the physics-informed Viterbi algorithm eliminates non-physical artifacts in the BrAC estimates.

[§]Corresponding Author; grosen@math.usc.edu.

1. Introduction

Transdermal alcohol biosensors (two such devices are shown in Fig. 1) use fuel cell technology to measure the amount of ethanol excreted through the skin after consuming alcoholic beverages. The devices provide near-continuous transdermal alcohol concentration readings, do not require active user participation, and can be worn for extended periods, making them potentially valuable to researchers interested in studying naturalistic alcohol use and clinicians treating alcohol use disorder. Researchers and clinicians have long relied on breath analyzers and self-report measures to monitor alcohol consumption outside of the laboratory/clinic. Breath analyzers require active participation by the user to obtain a reading at a single time point, and error can occur when readings are taken using improper breathing techniques, taken too close together in time, or taken too soon after drinking when alcohol remains in the mouth. Self-report methods such as drinking diaries can also be subject to error if a participant's memory is impaired by alcohol or if quantities of alcohol consumed are unknown in mixed drinks or in shared pitchers. Additionally, self-monitoring via breathalyzers and self-report are not a viable option if participants are motivated to conceal their alcohol use from researchers and clinicians. Thus, transdermal alcohol measurements provide an appealing alternative; they reduce user error and eliminate the need for costly, time-consuming in-person monitoring in the lab [29]. However, despite its many benefits, transdermal alcohol concentration (TAC) also has its limitations. Indeed, although generally positively correlated with blood and breath alcohol concentration (BAC and BrAC), accurately and consistently converting TAC into BAC or BrAC, the currently widely accepted measures of alcohol intoxication, poses a number of significant technical challenges. This has resulted in TAC sensors currently being used exclusively as abstinence monitors.

While BrAC is closely related to BAC [25], TAC is the result of the more complex physiological process of alcohol diffusing from the blood, through the skin, and into and measurement by, the sensor. This complexity leads to a number of confounding issues including 1) variation in transdermal alcohol concentration across people due to factors such as sex [19], age, body composition, and skin thickness, 2) across drinking episodes due to environmental factors such as ambient temperature and humidity, and alcohol dose [52], and 3) across devices due to differences in design, manufacturing, and calibration, all of which contribute to making estimation of BAC or BrAC via the forward or inverse filtering of TAC a challenging undertaking. To address this, our team has created an innovative new family of mathematical models that are trained a-priori, or off-line, using population data, and that produce quantitative estimates of BrAC from TAC along with conservative error bands.

Several researchers have attempted to create mathematical or computer models that could be used to convert TAC into BAC or BrAC with varying degrees of success. The problem of obtaining or estimating BrAC from a TAC biosensor signal in both earlier treatments and in our effort here is typically formulated as two serially connected inverse problems. In prior work, the first one, as illustrated in the upper left panel of Figure 2 involves collecting contemporaneous TAC data and either drinking diary or BrAC data from a cohort of subjects and then using it to fit (i.e. identify the unknown parameters in) a forward model. This forward model describes either the transport of ethanol from the blood through the skin

and its measurement by the sensor or from its ingestion through the mouth, to absorption into the blood in the gut, its metabolism in the liver, its transport through the skin and its measurement by the sensor. The second inverse problem, as illustrated in the upper right panel of Figure 2 and referred to as the prediction, inversion, or in the case of a linear model, deconvolution phase, involves using the fit forward model from the training phase and a subject's biosensor-measured transdermal alcohol concentration signal to obtain an estimate for the corresponding blood or breath alcohol concentration signal.

Some of the models developed for this purpose have taken the form of simple linear regression models ([12], [13], [19], [20]) that, for example, predict peak BrAC from TAC using explanatory variables which include sex, TAC peak, time to peak TAC, and area under the TAC curve. Other groups have looked at machine learning *black box* models. One approach involved using time series features of the TAC curve and random forests as predictors of BrAC [16], while our group has developed models based on physics-informed artificial neural networks and long short-term memory networks ([35], [36]). Until recently, the primary focus of our group has been first-principles physics based compartment (i.e. ODE) or distributed parameter (i.e. PDE) models or coupled hybrid systems involving both ([10], [15], [43]). Although these models were able to be used to deconvolve BrAC from TAC with reasonable accuracy, we found that they had to first be calibrated (i.e. have their parameters fit) to the individual participant wearing the sensor. The estimates they produced were also found to be sensitive to ambient environmental conditions. To mitigate these effects, we turned our attention to population-based models. This was achieved by allowing the parameters in the deterministic PDE models to be random with their distributions fit to population data. The objective was to eliminate the need for the participant to first complete a calibration drinking episode before wearing the biosensor and collecting data in the field ([18], [48], [49], [50], [51]).

In our treatment here, we develop a framework for estimating BrAC from biosensor measured observations of TAC based on what we shall refer to as a covariate-dependent, physics-informed, hidden Markov model (HMM). An HMM consists of an un-observable Markov chain together with observable but noisy output or emission variables. In our case here, the observations or emissions are BrAC and TAC. The hidden Markov states and their associated initial and transition probabilities represent the stochastic dynamics associated with the body's ingestion, transport, and metabolism of ethanol.

The basic idea of, and the theory behind, HMMs was established in a series of research papers by Baum and his colleges in the late 1960s and early 1970s [2, 3, 4, 5, 6], and combined into a single comprehensive tutorial in the paper by Rabiner [38]. Since then, they have been used across various applications as both a forward model and as a means of formulating and solving inverse problems. The most often cited application of HMMs dating back to the late 1980's is speech recognition [23, 38, 39, 45] with later applications following in the areas of acoustics, signal and image processing, genomics, and bioinformatics. They continue to enjoy widespread popularity with researchers in areas as diverse as musicology, gesture recognition, trend analysis, and finance [33]. A comprehensive bibliography of references for hidden Markov models and their application from 1989 through 2000 can be found in [7]. In addition, [8] is an excellent general resource

which, in addition to providing a rigorous treatment of the theory and derivations of the various computational algorithms for HMMs, includes an extensive survey of the literature. A more recent survey and review of the literature, essentially up to the present, covering the latest research on the theory and application of HMMs can be found in [33].

Our approach involves a number of new and innovative features including:

- Combining our prior knowledge and experience with neural networks and the first-principles modeling of the transdermal transport of ethanol, we have developed a way to augment and regularize a strictly data-driven HMM model with a physics-based PDE model with the subsequent effect of mitigating observed over-fitting phenomena.
- Our scheme solves the inverse problem of obtaining an estimate of BrAC from TAC by treating both quantities as HMM emissions and then using a hybrid of forward and inverse HMM filtering to obtain an estimate for BrAC, or eBrAC. This is in contrast to the typical HMM filtering problem wherein the inverse problem is solved via global decoding of the hidden Markov states.
- We are able to use local decoding of the hidden Markov states to quantify the uncertainty in our BrAC estimates and thereby provide conservative error bands along with the eBrAC signal.
- The underlying structure of our model is such that it offers the benefits of all the previous efforts described above to estimate BrAC from TAC. Indeed, 1) we are able to include the effects of un-modeled dependencies by allowing the HMM parameters to depend on covariates, 2) our model takes advantage of the demonstrated efficacy of a first-principles PDE model that captures the physics and physiology of the transdermal transport of ethanol, while at the same time benefits from the use of a data-driven, machine-learning based approach that mitigates the significant variability observed due to un-modeled effects, and 3) our model is a population model and thus does not require calibration to each individual subject, sensor and ambient environmental conditions.

The forward filtering problem for an HMM [28, 30] involves determining the most likely path (i.e. the sequence of states) through the hidden Markov chain based on a given sequence of noisy observations or emissions. The inverse filtering problem, on the other hand, involves determining the most likely model parameters (i.e. the Markov chain transition and initial state and the emission probabilities) given a set of sequences of observations or emissions. Given a fully trained HMM and a sequence of emissions, the Viterbi algorithm [54, 57], is used to solve the forward filtering problem, that is, to find the most likely path through the Markov chain. The Baum-Welch algorithm [57] is the standard method used to solve the inverse filtering problem. We note that both algorithms are considered to be unsupervised (since the hidden states are unobserved) learning algorithms based on maximum likelihood estimation and either the Expectation Maximization algorithm [9] in the case of Baum-Welch, or dynamic programming in the case of Viterbi. It is the underlying Markovian assumption on the hidden states that makes

the computations required to implement these algorithms combinatorially tractable. As these algorithms will appear frequently in the manuscript, we simply refer to them by their names.

We note that while the idea of combining physics-based modeling with machine learning techniques is relatively new, it is an area that has seen significant activity as of late (a comprehensive survey and extensive bibliography can be found in the recent review article [24]). However, virtually all of this work has been in the area of deep neural networks and deep learning (see, for example, [35, 40, 41, 42]). As we have noted earlier, while HMMs have been around for the past three to four decades, the idea of combining and regularizing them with physics-based models appears to be new. We are aware of only one instance of this in the literature in the form of a pre-print appearing in the archive [27].

A schematic or signal flow diagram illustrating our general approach is shown in the lower left and lower right panels of Figure 2 with the details of our scheme to follow in the subsequent sections below.

Remark In the application of interest to us here, modeling the ingestion, absorption, metabolism, and transport of ethanol in the human body as Markov chains can reasonably be motivated. Indeed, ethanol is highly miscible and finds its way, relatively rapidly, into all the water in the body. Consequently, in each time interval, the state of the hidden Markov chain is intended to capture 1) the subject's decision whether or not to take a drink, 2) the absorption of ethanol from the gut into the blood, 3) the mixing and distribution of the ethanol throughout the circulatory system and eventually the body water, 4) the enzyme (alcohol dehydrogenase) catalyzed oxidation of ethanol into a corresponding aldehyde and/or ketone in the liver, 5) the loss of ethanol contained in the body water through other mediums including urine, tears, breast milk, perspiration and exhalation (not including the portions measured by the sensors), etc. 6) the transport of ethanol from the blood and the body water into the dermal layer of the skin, 7) the exchange of ethanol molecules between the blood and body water and the gasses in the lung, 8) the transport of ethanol through the epidermal layer of the skin, and the flow of the gasses containing ethanol molecules into the collection chambers of either the TAC sensor or the breath analyzer. Processes 2, 3, 5, 6, 7, and 8 can all be described as diffusion which can be modeled via Brownian motion as a random walk which is of course a Markov chain (see, for example, [17]). Process 4 is typically modeled using Michaelis-Menten kinetics which can be shown to be well approximated to first order by an Euler scheme which can also be viewed as Markovian (see, for example, [10]). Finally, drinking patterns and behaviors have also been modeled with Markov chains and in fact HMMs (see, for example, [32], [47], and [55]).

An outline of the remainder of the paper is as follows. In Section 2 we review hidden Markov models and their associated algorithms for training and prediction, or more precisely, producing BrAC estimates based on biosensor measurements of TAC. In Section 3 we discuss the introduction of covariates into our model, and in Section 4 we discuss how our model can be used to quantify the uncertainty in the estimation through the local decoding of the hidden Markovian states. In Section 5 we propose the introduction of physics-informed regularization into our HMM that is based on a PDE model for the diffusion of ethanol through the epidermal layer of the skin. In Section 6 we describe how

our human subject data were collected and we present and discuss some of the results of our numerical studies. A final Section 7 has a summary and a few concluding remarks.

A note on notation: Since the various types of alcohol concentration and the HMM will be mentioned frequently throughout this paper, we have introduced and will continue to use the following acronyms: transdermal alcohol concentration will be referred to as **TAC**, breath alcohol concentration as **BrAC**, blood alcohol concentration as **BAC**, and hidden Markov model as **HMM**. To distinguish the actual signals from the estimated signals based on the model, the estimated signals have an “e” in front, that is **eTAC** and **eBrAC**. For improved readability, other than those mentioned here, in what follows, we will refrain from using any additional acronyms.

2. A Forward Full Body Alcohol Hidden Markov Model for Estimating BrAC from TAC

2.1. Hidden Markov Models

In this and the next section we review HMM notation, definitions, theory, and forward and inverse filtering algorithms. The details can be found in [8, 57].

An HMM is a statistical model that models transitions between different states of a process. The progression of states is assumed to be a Markov chain as described in [14, 17]. The states themselves and their transitions are unobservable, thus the term *hidden*. However, every state leads to an observable emission that can be used to gain insight into the Markov process. The architecture of an HMM is shown in Figure 3.

Let $\mathcal{O} = (o_1, \dots, o_T)$ be an observation sequence with d -dimensional observations $o_t = (o_{t,1}, \dots, o_{t,d})$ and let $\mathcal{X}$ be a random vector whose values are state sequences $\mathcal{X} = (X_1, \dots, X_T)$ with $X_i \in \{1, \dots, S\}$ for all i . Assume a p -dimensional vector of covariates $Z = (z_1, \dots, z_p)$. In classical Hidden Markov theory, the process $\mathcal{X}$ is a Markov process of order one [17], i.e.

$$P(X_{t+1} = j \mid X_1, \dots, X_t) = P(X_{t+1} = j \mid X_t).$$

Here we assume the Markov chain is time-homogeneous, that is

$$P(X_{t+1} = j \mid X_t = i) = P(X_2 = j \mid X_1 = i) \quad \forall t.$$

Typically the states of the Markov chain, being hidden, tend to be abstract and hence their interpretation relative to the application, requires some form of contextual decoding. More on this in Sections 4 and 6.8 below.

In the context of the alcohol biosensor problem, the observations are bivariate (i.e. $d = 2$) with a TAC and a BrAC component. Though not explicitly modeled, it is not unreasonable to consider the hidden states to be the generally unobservable discretized BAC level (and hence the terminology Hidden Markov Model; more on this later). Such an HMM is described by the parameters

- $\pi = (\pi_1, \dots, \pi_S)$, $\pi_i = P(X_1 = i)$, i.e. the prior probabilities of the chain $\mathcal{X}$ to start in state i ,
- $a_{i,j} = P(X_{t+1} = j | X_t = i)$ with $a_{i,j} \geq 0$ and $\sum_{j=1}^S a_{i,j} = 1$. These parameters describe the transition probabilities between different states.
- $b_i(o_t) = P(o_t | X_t = i)$, the probability of observing o_t given that the current state is i . These are called emissions. For the density functions b_i various choices, both discrete and continuous, are possible. Here, we will assume the emission distribution to be Gaussian; more on this to follow.

In the sequel we combine all the parameters into a single parameter vector, θ .

2.2 Training and Initialization

Training an HMM involves determining the parameters θ based on sequences of observations $\mathcal{O}$. We consider the following quantities that can be computed efficiently by using the Baum-Welch or forward-backward algorithm [38]. Briefly, as in [38], we denote the probability of the partial observation sequence $(o_1, \dots, o_{t-1})$ and being in state i at time t given the model parameters θ by $\alpha_t(i) = P(o_1, \dots, o_{t-1}, X_t = i | \theta)$, the probability of the partial observation sequence $(o_t, \dots, o_T)$ given that the state sequence is in state i at time t by $\beta_t(i) = P(o_t, o_{t+1}, \dots, o_T | X_t = i, \theta)$. With that, the probability of going from state i to state j given the observations $\mathcal{O}$ can be expressed by (see [38])

$$\gamma_t(i, j) = P(X_t = i, X_{t+1} = j | \mathcal{O}, \theta) = \frac{\alpha_t(i) a_{i,j} b_j(o_{t+1}) \beta_{t+1}(j)}{\sum_{m=1}^S \alpha_t(m) \beta_{t+1}(m)},$$

using the definition of conditional probability and the law of total probability. Further, the probability of being in state i at time t given the observations $\mathcal{O}$ is then given by

$$\gamma_t(i) = P(X_t = i | \mathcal{O}, \theta) = \sum_{j=1}^S P(X_t = i, X_{t+1} = j | \mathcal{O}, \theta) = \sum_{j=1}^S \gamma_t(i, j). \quad (1)$$

We note that all of these quantities can be efficiently computed recursively [38].

The optimal choice for the parameters, $\theta = \theta^*$, is determined using the Baum-Welch algorithm that is a special case of an expectation maximization algorithm [9]. In general, expectation maximization algorithms provide an iterative method to find maximum likelihood estimates of parameters in statistical models that contain unobserved variables [11]. As the state transitions in HMMs are hidden, such an algorithm is the natural choice for the parameter estimation. The log likelihood ℓ of the observations with parameter θ is given by $\ell(\theta | \mathcal{O}, \mathcal{X}) = \log \mathcal{L}(\theta | \mathcal{O}, \mathcal{X}) = \log P(\mathcal{O}, \mathcal{X} | \theta)$, where $\mathcal{L}$ denotes the likelihood. As is standard in the literature of expectation maximization [11], we define

$$Q(\theta, \theta^{(k-1)}) = \mathbb{E}_{\theta^{(k-1)}}(\ell(\theta | \mathcal{O}, \mathcal{X})) = \sum_{\mathcal{X}} \log(P(\mathcal{O}, \mathcal{X} | \theta)) P(\mathcal{O}, \mathcal{X} | \theta^{(k-1)}) \quad (2)$$

where $\theta^{(k-1)}$ denotes the $(k-1)^{\text{st}}$ iterate of θ . The joint probability of an observation sequence $(o_1, \dots, o_T)$ and a state sequence $(X_1, \dots, X_T)$ can be written as the product of: the

initial probability to start in state X_1 , the probability of observing o_1 given the state X_1 , and the product of the transition probabilities from X_{t-1} to X_t and observing o_t given the state X_t for $t = 2, \dots, T$. That is,

$$P(\mathcal{O}, \mathcal{X} \mid \theta) = \pi_{X_1} b_{X_1}(o_1) \prod_{t=2}^T a_{X_{t-1}, X_t} b_{X_t}(o_t),$$

and therefore (2) turns into

$$\begin{aligned} Q(\theta, \theta^{(k-1)}) &= \sum_{\mathcal{X}} \log \left(\pi_{X_1} b_{X_1}(o_1) \prod_{t=2}^T a_{X_{t-1}, X_t} b_{X_t}(o_t) \right) P(\mathcal{O}, \mathcal{X} \mid \theta^{(k-1)}) \\ &= \sum_{\mathcal{X}} \log(\pi_{X_1}) P(\mathcal{O}, \mathcal{X} \mid \theta^{(k-1)}) + \sum_{\mathcal{X}} \left(\sum_{t=2}^T \log(a_{X_{t-1}, X_t}) \right) P(\mathcal{O}, \mathcal{X} \mid \theta^{(k-1)}) \\ &\quad + \sum_{\mathcal{X}} \left(\sum_{t=1}^T \log(b_{X_t}(o_t)) \right) P(\mathcal{O}, \mathcal{X} \mid \theta^{(k-1)}). \end{aligned} \quad (3)$$

This expression for Q is maximized by maximizing the three terms individually. For the first term involving the prior probabilities $\pi_{\hat{p}}$ partitioning all sequences, $\mathcal{X}$, by their initial state, we get

$$\sum_{\mathcal{X}} \log(\pi_{X_1}) P(\mathcal{O}, \mathcal{X} \mid \theta^{(k-1)}) = \sum_{i=1}^S \log(\pi_i) P(\mathcal{O}, X_1 = i \mid \theta^{(k-1)}). \quad (4)$$

Maximizing (4) with respect to θ with the constraint $\sum_{i=1}^S \pi_i = 1$ yields

$$\pi_i^* = \frac{P(\mathcal{O}, X_1 = i \mid \theta^{(k-1)})}{P(\mathcal{O} \mid \theta^{(k-1)})} = \gamma_1(i).$$

For the second term, partitioning all sequences $\mathcal{X}$ by their states X_{t-1} and X_t we find that

$$\begin{aligned} &\sum_{\mathcal{X}} \left(\sum_{t=2}^T \log(a_{X_{t-1}, X_t}) \right) P(\mathcal{O}, \mathcal{X} \mid \theta^{(k-1)}) \\ &= \sum_{i=1}^S \sum_{j=1}^S \sum_{t=2}^T \log(a_{i,j}) P(\mathcal{O}, X_{t-1} = i, X_t = j \mid \theta^{(k-1)}). \end{aligned}$$

For this term, maximizing under the constraint that $\sum_{j=1}^S a_{i,j} = 1$ yields

$$a_{i,j}^* = \frac{\sum_{t=1}^T P(\mathcal{O}, X_{t-1} = i, X_t = j \mid \theta^{(k-1)})}{\sum_{t=1}^T P(\mathcal{O}, X_{t-1} = i \mid \theta^{(k-1)})} = \frac{\sum_{t=1}^T \gamma_{t-1}(i, j)}{\sum_{t=1}^T \gamma_{t-1}(i)}. \quad (5)$$

We will return to the maximization of the third term in (3) in the next section when we discuss the incorporation of covariates in our emissions model.

In our HMM the emissions are bivariate (i.e. BrAC and TAC). Consequently we fit a bivariate Gaussian for each state. The density is given by

$$b_i(y) = f(y | X = i) = \frac{1}{(2\pi)^{d/2} |\Sigma_i|^{1/2}} \exp\left(-\frac{1}{2}(y - \mu_i)^T \Sigma_i^{-1} (y - \mu_i)\right), \quad (6)$$

where in (6) $y \in \mathbb{R}^2$ denotes a possible value for an emission random vector o_t , μ_i denotes a mean vector, and Σ_i denotes a covariance matrix. Since the emissions are known to be non-negative, another reasonable choice here would be, for example, a bivariate log-normal distribution. However, we opted for a normal to take advantage of the ease of parameterization and the fact that the marginal distributions which are of particular significance to us here are normal as well. We rely on the training data to produce means, μ_i , and covariance matrices, Σ_i , that make negative values highly unlikely. The assumption that the observations of BrAC and TAC have a bivariate normal distribution is essentially the standard statistical model that BrAC and TAC at each sample time take values that are functions of the current state of the hidden Markov chain X_t plus zero mean, (in general) correlated Gaussian noise with covariance also a function of the current state of the hidden Markov chain. That is $Y_t = y(X_t) + \epsilon_t$, where $\epsilon_t \sim \mathcal{N}(0, \Sigma(X_t))$ and $\mathcal{N}(\mu, \Sigma)$ denotes the normal distribution with mean μ and covariance matrix Σ .

Training an HMM is a non-linear optimization problem over a large set of parameters. The Baum-Welch algorithm is likely to find local maxima. Hence the initial parameter guesses θ^0 are of great importance. While there are initialization methods for the transition probabilities [26], we focus on initialization of the means μ_i and covariance matrices Σ_i of the emission probabilities.

One possible initialization of the means μ_i^0 makes use of a common clustering algorithm, the k -means algorithm [31]. Given a number of clusters k , here the number of different states $k = S$, this algorithm minimizes the error function

$$\sum_{i=1}^k \sum_{o_j \in C_i} \|o_j - \mu_i^0\|^2$$

where C_i describes the i -th cluster, μ_i^0 is the mean or center of this cluster, and o_j is the j -th observation. In this way, the observations are clustered into a number of different levels of BrAC and TAC that correspond to the number of hidden states S . Using this initialization for the means, the hidden states will tend to emit different TAC and BrAC levels. Then, having estimated values for μ_i , it is straightforward to produce estimates for the covariance matrices Σ_i . For each cluster $i = 1, \dots, S$, we can estimate the initial covariance matrix as $\Sigma_i^0 = \frac{1}{|C_i|} \sum_{o_j \in C_i} (o_j - \mu_i^0)(o_j - \mu_i^0)^T$. In the covariate-dependent emissions case, which is of particular interest to us here and will be discussed in the next section, this initialization process still applies by setting the regression coefficients to zero for all states.

3. Introducing Covariates and Using the HMM to Solve the Inverse Problem of Obtaining eBrAC from TAC

In the alcohol biosensor problem, it is natural to assume that the process of ethanol diffusing through the epidermal layer of the skin depends on covariates such as body features or ambient conditions. Consequently, it is reasonable to model the transitions between various states or the emission distributions as depending on these covariates.

3.1. Covariate-Dependent State Transitions

Making state transitions covariate dependent is straightforward [57] and it often reveals additional insight [47]. In order to guarantee that the matrix of transition probabilities for the hidden Markov chain is a stochastic matrix, i.e. the transition probabilities sum to 1, the standard approach for introducing covariates into the transition probabilities is to use multinomial logistic regression. In this case, as in [57], we set

$$a_{i,j}^{\tilde{Z}} = P(X_t = j \mid X_{t-1} = i, \tilde{Z}) = \frac{\exp(\phi_{i,j}^T \tilde{Z})}{\sum_{j=1}^S \exp(\phi_{i,j}^T \tilde{Z})}, \quad (7)$$

where, so as to include a bias term when identifying the model, $\tilde{Z}$ denotes an augmented vector of covariates $\tilde{Z} = (1, Z)$, and $\phi_{i,j}$ denote the weight vectors to be determined. We imposed the additional restriction that $\phi_{i,i} = 0$ for all states i . The model (7) maintains the assumption of a time-homogeneous Markov process. It is possible to model transition probabilities $a_{i,j}^{\tilde{Z}_t}$ that depend on a time-dependent covariate vector Z_t . The resulting Markov chain would not be homogeneous. The covariates of interest to us here are constant with respect to time and hence our model with covariates will continue to be homogeneous with respect to time.

When using the model (7) for the transition probabilities, (5) is no longer valid. In fact, a closed form expression of the corresponding maximizer $a_{i,j}^{\tilde{Z}_t*}$ does not exist. Consequently a numerical optimization technique is required.

3.2. Covariate-Dependent Emissions or Observations

In our model we also allowed the HMM emission distributions to be dependent on the covariates. To include covariates in the emissions model, we replace the density given in (6) with a bivariate Gaussian distribution whose mean is a linear function of the covariates. That is, a Gaussian with an added regression term for its mean to account for the covariates. The standard density of such a distribution [9] is given by

$$b_i(y) = f(y \mid Z, X = i) = \frac{1}{(2\pi)^{d/2} |\Sigma_i|^{1/2}} \exp\left(-\frac{1}{2}(y - B_i Z - \mu_i)^T \Sigma_i^{-1} (y - B_i Z - \mu_i)\right), \quad (8)$$

where μ_i denotes a nominal mean vector, Σ_i denotes a covariance matrix, and B_i denotes a regression matrix for the covariates. Once again, by augmenting Z and B as $\tilde{Z} = (1, Z)$ and $\tilde{B}_i = (e, B_i)$, respectively, with $e = (1, \dots, 1)^T$, we get $y - B_i Z - \mu_i = y - \tilde{B}_i \tilde{Z}$, i.e. that μ_i in the

above formulation is a covariate-independent bias for the mean of $y|(X = i)$. We again note that a time-dependent covariate vector Z_t could also be used here.

Now, returning to the maximization of the third term in (3), the one dealing with the state-dependent emission probabilities, $b_i(o_t)$, we write

$$\begin{aligned} & \sum_{\mathcal{X}} \left(\sum_{t=1}^T \log(b_{x_t}(o_t)) \right) P(\mathcal{O}, \mathcal{X} \mid \theta^{(k-1)}) \\ &= \sum_{i=1}^S \sum_{t=1}^T \log(b_i(o_t)) P(\mathcal{O}, X_t = i \mid \theta^{(k-1)}). \end{aligned} \quad (9)$$

In [34], it is shown how the maximum likelihood estimators for the conditional linear Gaussian distribution can be computed. In the case of $B_i = 0$ for all i , i.e. when there are no regression terms, from (9) the maximum likelihood estimator for μ_b, μ_i^* is obtained as

$$\mu_i^* = \frac{\sum_{t=1}^T P(X_t = i \mid \mathcal{O}, \theta) o_t}{\sum_{t=1}^T P(X_t = i \mid \mathcal{O}, \theta)} = \frac{\sum_{t=1}^T \gamma_t(i) o_t}{\sum_{t=1}^T \gamma_t(i)}$$

with $\gamma_t(i)$ as in (1). In the case $S = 1$, this agrees with the sample mean.

In the covariate regression case, μ_i and B_i have to be estimated simultaneously. Consequently the maximum likelihood estimator for the augmented matrix $\tilde{B}, \tilde{B}^*$, is given by the solution of a weighted version of the normal equations as

$$\tilde{B}_i^* = \left(\sum_{t=1}^T \gamma_t(i) o_t \tilde{Z}^T \right) \left(\sum_{t=1}^T \gamma_t(i) \tilde{Z} \tilde{Z}^T \right)^{-1}. \quad (10)$$

It then remains to find an estimate of the covariance matrices Σ_b, Σ_i^* . As shown in [34], using the maximum likelihood estimator $\tilde{B}_i^*$ from (10) this is found to be

$$\Sigma_i^* = \frac{1}{\sum_{t=1}^T \gamma_t(i)} \left(\sum_{t=1}^T \gamma_t(i) o_t o_t^T - \tilde{B}_i^* \sum_{t=1}^T \gamma_t(i) \tilde{Z} o_t^T \right).$$

3.3. Using the HMM to Estimate or Predict BrAC from Observations of TAC

After training the model with the observed BrAC and TAC data, we want to use the HMM to produce, or predict, an estimated BrAC signal, eBrAC, based only on a TAC signal. During the training, we provide both TAC and BrAC, so, as noted in the previous section, we fit either a covariate-independent (i.e. (6)) or a covariate-dependent (i.e. (8)) bivariate Gaussian for every state.

In the prediction stage, we use the marginal distribution of $b_i(y_1, y_2)$ with respect to the first variable y_1 , that is, the TAC observation. The marginal density takes the form

$$b_i(y_i) = f(y_i | Z, X = i) = \frac{1}{\sqrt{2\pi\Sigma_{i1}}} \exp\left(-\frac{1}{2\Sigma_{i1}}(y_i - B_i^1 Z - \mu_{i1})^2\right), \quad (11)$$

where B_i^1 denotes the first row of B_i , μ_{i1} is the first entry of μ_i and Σ_{i1} is the variance of $y_1|(X = i)$. Using the TAC observations and the conditional densities $b_i(y_1)$ given in (11), we use the Viterbi algorithm [54] to predict the most likely state sequence based on the observation sequence $\mathcal{O}_1 = (o_{1,1}, \dots, o_{T,1})$. Then, given the most likely (or, for that matter, any) state sequence $\mathcal{X} = (X_1, \dots, X_T)$, it is then straightforward to evaluate the eBrAC signal, $\tilde{\mathcal{O}}_2$, as $\tilde{\mathcal{O}}_2 = (\mu_{x_{12}}, \dots, \mu_{x_{T2}})$, where $\mu_{x_{12}}$ denotes the second entry of μ_{X_1} i.e. the mean BrAC value that is associated with the state X_1 . Note that by providing the Viterbi algorithm with only the marginal densities $b_i(y_1)$, there is some inherent loss of information. However, the training of the HMM with TAC and BrAC should enable the model to yield meaningful predictions even when only observations of TAC are provided. In Section 5 we outline a physics-informed regularization for this process that can help to reduce the impact of this loss of information.

3.4. Model Complexity

When introducing covariates into the HMM model, it is important to consider the induced changes to the model complexity. In the classical case, i.e. a multivariate Gaussian distribution without a regression on the covariates for the emissions and S different states with covariate-independent transitions, we need to estimate $S \cdot d$ parameters for the means μ_i , $S \cdot d(d+1)/2$ parameters for the positive definite symmetric covariance matrices Σ_i , S priors and $S(S-1)$ transition probabilities. Remember that d denotes the dimension of the observations. Now, in the case of covariate-dependent state transitions as in (7), instead of $S(S-1)$ transition probabilities, we need to estimate the vectors $\phi_{i,j}$. We assumed that $\phi_{i,i} = 0$, and so there are $S(S-1) \cdot p$ parameters to be estimated for the transitions, where p again denotes the number of covariates. In the case of covariate-dependent emissions, we add $S \cdot d \cdot p$ parameters for the regression matrices B_i . Hence, it can clearly be seen that the dependence on covariates leads to significantly increased model complexity, especially in the case of covariate-dependent transitions. To accurately fit a model with large complexity typically requires a large training data set. In the context of the alcohol biosensor problem, data collection with human subjects is an expensive and time consuming undertaking, and so available data sets are relatively small. Consequently, including a large number of covariates does not necessarily lead to a model with enhanced predictive power. Instead, one should balance the greater flexibility of a more complex model with the scarcity of training data. This can be achieved by including only a small number of highly significant covariates, as we discuss in Section 6.

4. Decoding the Hidden States and Uncertainty Quantification

Since HMMs are probabilistic models, the most likely path estimation and hence the BrAC estimation are subject to statistical considerations. The probabilistic formulation of an HMM can be utilized to gain insight into the uncertainty that comes with the estimation. In this

section, we describe how we can quantify the uncertainty and how this leads to credible bands for eBrAC.

As described in Section 3.3, the Viterbi algorithm can be used to find the most likely state sequence $\mathcal{X}^* = (i_1^*, \dots, i_T^*)$ given the observations $\mathcal{O}$ as $\mathcal{X}^* = \operatorname{argmax}_{\mathcal{X}} P(\mathcal{X} | \mathcal{O}, \theta)$. This is referred to as *global decoding* [57] and the result is guaranteed to be a feasible state sequence in the sense that $P(X_{t+1} = i_{t+1}^* | X_t = i_t^*) > 0$.

Other than the global decoding that considers the full sequence of states, the process of *local decoding* finds the most likely state i_t^* at every time step t ; that is $i_t^* = \operatorname{arg max}_{i \in \{1, \dots, S\}} P(X_t = i | \mathcal{O}, \theta) = \operatorname{arg max}_{i \in \{1, \dots, S\}} \gamma_t(i)$ with $\gamma_t(i)$ as in (1). Note that local decoding is a byproduct of the forward-backward algorithm since it is based on the quantities $\gamma_t(i)$. The main disadvantage of local decoding is that it potentially can yield an unfeasible or impossible state sequence. However, local decoding not only leads to an identification of the most likely state at time t , but it also provides a measure of confidence in the form of $P(X_t = i_t^*) = \gamma_t(i_t^*)$. This probability can be used to inform the uncertainty that comes with the state decoding. More on this in Section 6 below.

A second source for uncertainty are the emissions. In our model, we fit bivariate conditional linear Gaussians that are parameterized by a covariate-independent mean μ , a regression matrix B and a covariance matrix Σ . Hence, given a state i , the emission o is subject to the uncertainty introduced by the variance Σ_i .

Thus, there are two types of uncertainty: A state decoding-related one and an emission-related one. Fortunately, both uncertainties can be quantified by parameters of the fitted model. In Section 6 below we present a method that takes into account both types of uncertainty to yield credible bands for the eBrAC signal.

5. A Physics-Informed Viterbi Algorithm

In this section, we extend the HMM-based scheme for estimating BrAC from TAC to one that combines the data-driven nature of the HMM approach described previously with prior knowledge in the form of a first principles physics-based model given by a hybrid system of partial and ordinary differential equations. This combined approach in the context of artificial neural networks has been shown to yield improved accuracy in situations where sufficient training data may be difficult, costly, or impossible to obtain and thus hindering the learning process in a purely data-driven regime (see, for example, [35], [40], [41], [42], [46], and [56]). We show that such an approach is also well-suited for the alcohol biosensor problem and our HMM-based solution as well.

It has been shown (see, for example, [15] and [43]) that, with reasonable accuracy, the following first principles physics-based hybrid input/output system of partial and ordinary differential equations models the transport of ethanol from the blood through the epidermal layer of the skin and into the collection chamber of the TAC sensor.

Let $\Lambda > 0$ denote the thickness of the epidermal layer of the skin in centimeters and let η be the depth in centimeters in the epidermal layer, $0 \leq \eta \leq \Lambda$. $\eta = 0$ denotes the skin surface and

$\eta = \Lambda$ denotes the boundary between the epidermal and dermal layers. We let t denote time in hours, and consider the system

$$\begin{aligned} \frac{\partial x}{\partial t}(t, \eta) &= \alpha \frac{\partial^2 x}{\partial \eta^2}(t, \eta), \quad 0 < \eta < \Lambda, \quad 0 < t < T, \\ \frac{d\tilde{w}}{dt}(t) &= \frac{\gamma\alpha}{\rho_{P/A}} \frac{\partial x}{\partial \eta}(t, 0) - \delta \tilde{w}(t), \quad 0 < t < T \\ x(t, 0) &= \rho_{P/A} \tilde{w}(t), \quad \alpha \frac{\partial x}{\partial \eta}(t, \Lambda) = \beta \rho_{P/A} u(t), \quad 0 < t < T, \\ \tilde{w}(0) &= \tilde{w}_0, \quad x(0, \eta) = \varphi_0(\eta), \quad 0 < \eta < \Lambda, \\ y(t) &= \theta \tilde{w}(t), \quad 0 < t < T. \end{aligned} \quad (12)$$

The state $x(t, \eta)$, is the concentration of ethanol in the interstitial fluid at time t and depth η in the epidermal layer, the state $\tilde{w}(t)$ is the concentration of ethanol in the air in the TAC sensor collection chamber at time t , the input $u(t)$ is the BrAC at time t , and the output $y(t)$ is the TAC at time t , where all units are in grams per milliliter. The duration of the drinking episode in hours is denoted by T . $\tilde{w}_0$ and φ_0 are the initial conditions for $\tilde{w}$ and x , respectively, and we have assumed that there is no ethanol in either the epidermal layer or the TAC biosensor collection chamber at time $t = 0$, that is $\tilde{w}_0 = 0$ and $\varphi_0 = 0$. The parameter $\alpha > 0$ denotes the effective diffusivity of ethanol in the interstitial fluid in the epidermal layer, $\beta > 0$ denotes the effective linear flow rate at which capillary blood plasma from the dermal layer fills up the interstitial fluid in the epidermal layer, and $\rho_{P/A}$ denotes the Henry's law (see, for example, [21], [25]) partition coefficient for ethanol in blood plasma and air at normal body temperature (see [21] and [22]). In the model for the TAC collection chamber, it is assumed that the inflow of ethanol is proportional to the flux out of the epidermal layer at the surface of the skin, $\alpha \frac{\partial x}{\partial \eta}(t, 0)$, with proportionality constant γ , and the outflow is simply proportional to the concentration of ethanol in the collection chamber with proportionality constant δ . The output gain, θ , is the factor the TAC sensor uses to convert the concentration of ethanol in the collection chamber into a TAC.

Since the only observed and observable quantities in (12) are BrAC input, u , and TAC output, y , with the state variables x and w remaining *hidden*, elementary changes of variable (see, for example, [58]) allow for the transformation of the system (12) into an equivalent system involving only four parameters which we shall denote by the vector $q = [q_1, q_2, q_3, q_4]$ in the positive orthant of $\mathbb{R}^4$. Moreover, since the observed output y is time-sampled (with sampling time τ dictated by the sensor hardware) and if we assume the input u is zero-order hold, using standard techniques from the theory of linear semigroups for abstract parabolic systems [37, 53] together with spline based Galerkin approximations (for which there exists a rigorous convergence theory [18]), the system (12) can be equivalently written as a single input-single output convolution system as

$$y = h(q) * u \quad \text{or} \quad y_k = \sum_{j=0}^{k-1} h_{k-1-j}(q) u_j, \quad (13)$$

where the q -dependent convolution kernel, or filter, $h(q) = \{h(q)_k\}$ is given by $h_k = C(q)e^{k\tau A(q)}B(q) = C(q)(e^{\tau A(q)})^k B(q)$ with the matrices $C(q)$, $A(q)$, and $B(q)$ all given in

closed form as functions of the parameters q (see [18]). The parameters, q_i , $i = 1, 2, 3, 4$ appearing in (13) are functions of the physical parameters in (12). While some of these physical parameters can be found in the literature (typically determined empirically in-vitro in the laboratory), we instead estimated the q_i 's statistically by first fitting the model (13) to data from individual drinking episodes [10],[43] and then averaging those estimates over the entire population of episodes to obtain a crude PDE-based population model. Without too much additional effort, it is possible to obtain a more sophisticated population model (see, for example, [1], [18], [50], [51], and [58]), but since this study is still at the proof-of-concept stage, we did not do this.

There are some choices as to how the training and subsequent use of our BrAC estimating HMM model could be augmented with physics-informed regularization based on the model given in (13). However, introducing physics-based information into the HMM training stage is problematic. The key to the success of HMMs is the exploitation of the Markovian assumption which allows for the highly efficient implementation of the Expectation Maximization algorithm. This is based on the observation that the step in the algorithm in which the expectation, or the quantity Q given in (3), is maximized, nicely decomposes into the maximization of three terms involving the prior probabilities, the transition probabilities, and the distribution of the emissions. It seems unavoidable that if regularization were to be introduced here, the elegant and efficient structure of the Baum-Welch algorithm will be disrupted resulting in significant degradation of the training algorithm's performance.

An alternative, on the other hand, is to introduce physics-informed regularization in the prediction or BrAC estimation stage. Once the HMM has been trained using observed BrAC and TAC data, we use the HMM to estimate the BrAC signal corresponding to a given TAC signal. As part of the training stage, we fit bivariate Gaussians (that depend on the state of the hidden Markov chain) to the emissions or outputs, $o_t = [o_{t,TAC}, o_{t,BrAC}]$. We then use the given TAC signal, their corresponding conditional densities, and the Viterbi algorithm to predict the most likely path of states through the Markov chain. We obtain the estimated BrAC from the appropriate Gaussian marginal distribution. (Recall that we observed that although providing the Viterbi algorithm with only the marginal TAC densities is an inherent loss of information, since the training of the HMM used both BrAC and TAC, the resulting forward model yields meaningful estimates of BrAC even when only observations of TAC are provided.) It is at this point that prior physics-based insight can be introduced into the Viterbi algorithm, creating a physics-informed Viterbi algorithm which will yield a path through the hidden states that is consistent with both the data and the first principles physics-based model given by (12) and (13).

To make this precise, we recall that the two key quantities in the classical Viterbi algorithm are the arrays ξ and ψ , where $\xi(j, t)$ is the probability of the most likely path $(X_1, X_2, \dots, X_t)$ with $X_t = j$, $j = 1, 2, \dots, S$, based on the observed emissions sequence $(o_1, o_2, \dots, o_t)$, and $\psi(j, t)$ stores the state X_{t-1} of the most likely path $(X_1, X_2, \dots, X_{t-1}, X_t = j)$, $j = 1, 2, \dots, S$.

Initially, when $t = 1$, for $j = 1, 2, \dots, S$, we set $\xi(j, 1) = \pi_j b_j(o_1)$, where b_j is given by (6), and $\psi(j, 1) = 0$ since no predecessor state exists. Then, by sequentially filling the arrays ξ and ψ , the most likely path is computed. In our physics informed Viterbi algorithm, we

introduce the array $\kappa(j, t)$ that stores the predicted TAC value at time t , based on the most likely path $(X_1, X_2, \dots, X_{t-1}, X_t = j)$, $j = 1, 2, \dots, S$. To compute the values for κ , for $j = 1, 2, \dots, S$, we use the BrAC sequence $\tilde{u}$ derived from the most likely path $(X_1, X_2, \dots, X_{t-1}, X_t = j)$ using the mean values of the state-dependent emission distributions as described in Section 3.3. We compute the corresponding physics-based most likely TAC path $\tilde{y}$ using the model (13) as $\tilde{y} = \mathcal{L}(\tilde{u}) = h * \tilde{u}$ and set $\kappa(j, t) = \tilde{y}_t$, $j = 1, 2, \dots, S$. We then augment the problem of finding the next most likely state with a weighting of the absolute difference between y_{t-2} and $\kappa(\cdot, t-2)$. In this way, the objective of the Viterbi algorithm is no longer to simply find the most likely path through the Markov chain based only on the probabilistic formulation of the HMM, but rather to find a path that best matches both the fitted HMM model and the physics-based model. The precise algorithm is shown as a pseudocode. In Section 6, we provide numerical results using this approach and compare them to results that were obtained without this physics-regularization.

Algorithm 1 Physics-Informed Viterbi

```

function PHYSVITERBI
  for  $j = 1 : S$  do
     $\xi(j, 1) = \pi_j \cdot b_j(y(1))$ 
     $\psi(j, 1) = 0$ 
  end for
  for  $t = 2 : T$  do
    for  $j = 1 : S$  do
       $\tilde{X}_{t-1} = j$ 
      for  $k = t - 2, \dots, 1$  do
         $\tilde{X}_k = \psi(\tilde{X}_{k+1}, k + 1)$ 
      end for
       $\tilde{u} = (\mu_{\tilde{X}_{1_2}}, \dots, \mu_{\tilde{X}_{t-1_2}})$ 
       $\tilde{y} = \mathcal{L}(\tilde{u}) = h * \tilde{u}$ 
       $\kappa(j, t - 1) = \tilde{y}(t - 1)$ 
       $\psi(j, t) = \operatorname{argmax}_i \left( \xi(i, t - 1) \cdot a_{i,j} \right. \\ \left. - \beta |y(t - 2) - \kappa(i, t - 2)| \right)$ 
       $\xi(j, t) = \xi(\psi(j, t), t - 1) \cdot a_{\psi(j, t), j} \cdot b_j(y(t))$ 
    end for
  end for
end function

```

6. Numerical Studies

6.1. BrAC, TAC, and Covariate Data

The data consists of BrAC, TAC, and covariate data from a sample of 40 participants (50% female, ages 21 – 33 years, 35% BMI ≥ 25.0), who each completed four laboratory drinking sessions. Each participant consumed the same amount of alcohol in all four sessions, which was designed to reach a maximum BrAC peak of .080 (the legal limit for intoxication in the US) for that person based on body water weight. There were three patterns of consumption: participants consumed the alcohol in 15 minutes (Single Dose) in two of the sessions, in two 15-minute periods spaced 30 minutes apart (Dual Dose) in one session, and over 60 minutes (Steady Dose) in one session. An Intoximeter IV breath analyzer (Intoximeters, Inc., St. Louis, MO) was used to obtain BrAC, and two SCRAM-CAM devices (Alcohol Monitoring

Systems, AMS, Littleton, CO) were placed on the upper arms to obtain TAC. Measurements were taken at 10-minute intervals over the entire BrAC curve. After BrAC returned to .000, the TAC devices continued to record automatically at 30-minute intervals until TAC also returned to .000. Several sets of BrAC-TAC data were excluded from analyses due to a change in the protocol, and with two sets of TAC data recorded in each session, our final sample consisted of 256 BrAC-TAC datasets [44]. We also obtained covariates that may affect the BrAC-TAC relationship, including person-level covariates of sex, age, race/ethnicity, body measurements (e.g. hip and waist measurements, percent body fat), and past 90-day drinking behavior and session-level covariates of alcohol dosing pattern and blood pressure readings (baseline, minimum, and maximum).

6.2. Numerical Results

We split the 256 sets of BrAC-TAC data into a training set containing 150 episodes and a test set with the remaining 106 episodes, with all of each participant's episodes placed either in the training or in the test set to reduce the effects of non-independence of repeated drinking episodes within participants. All results displayed below show estimates on the test set using a model that has been trained solely on the training set. The covariate-dependent HMM was implemented in MATLAB based on the Hidden Markov Model (HMM) Toolbox. Through experimentation on the number of hidden states in our model, we settled on $S = 20$ hidden states because this number appeared to yield the best representation of the data. The particular episodes used to demonstrate our approach were selected because they were generally representative of the particular aspect of our method we were intending to illustrate.

6.3. An HMM without Covariates

In Figure 4, the upper row shows TAC and BrAC data together with their respective predictions, eBrAC and eTAC, without including covariates when both TAC and BrAC are inputs to the Viterbi algorithm. The lower row shows eBrAC predictions based on the TAC signal without including BrAC or covariates. In the bivariate case, the predicted signals provide accurate approximations of the true data, especially for the BrAC data. This is not surprising in so far as the BrAC signals in the data set tend to show less variation than do the TAC signals. Overall we observe that the mechanism of using the observations to find the most likely state sequence and then predicting the signal by means of the corresponding emissions works quite well. That is, the trained HMM appears to appropriately model the data and capture the underlying dynamics.

For the case of BrAC prediction based only on observations of TAC, it is not surprising that we see that the accuracy decreases somewhat. Nevertheless, the eBrAC signal still shows remarkably good agreement with the true signal with regard to its shape and amplitude. Table 1 shows the median of the relative errors in the eBrAC and eTAC signals, peak eBrAC and eTAC values, and the area underneath the eBrAC and eTAC curves (AUC) where eBrAC and eTAC were estimated using the Viterbi algorithm and either BrAC and TAC or TAC alone. The table also shows the absolute error (in units of hours) for the time at which the peak eBrAC and eTAC values occurred. Errors for eTAC are lower when the Viterbi algorithm is supplied with only TAC and the eBrAC errors are much lower when both TAC

and BrAC are used as the basis for prediction. We observed no clear difference between errors for the test set and for the training set. From this we infer that our model generalizes well with no evidence of over-fitting. With regard to how the model would perform when used in practice to estimate BrAC based on TAC, we observed that the median relative peak error was just over 20%, the median relative AUC error was just over 25% and the predicted peak time deviated from the true peak time by about 15 minutes. By way of a comparison, the results in Table 1 are consistent with, but somewhat less accurate than, similar results obtained using a nonlinear least squares, single subject, single laboratory alcohol challenge episode trained deterministic diffusion equation model that was then tested on that same subject's 10 additional drinking episodes conducted in the field [10, 15, 43]. We note that this is a much smaller and clearly significantly less diverse data set than the one used to train and test our HMM models here.

6.4. A Covariate-Dependent HMM

To select the covariates that are most significant for determining eBrAC from TAC to include in our HMM model, we introduce the scalar quantity

$$\lambda = \frac{\max_{t \in \{1, \dots, T\}} O_{t,2}}{\max_{s \in \{1, \dots, T\}} O_{s,1}}, \quad (14)$$

which is the ratio of the maximum BrAC value and the maximum TAC value during a drinking session. We then use this scalar quantity as the dependent or response variable in a standard linear regression using all available covariates as explanatory variables. Next we use the resulting fit statistics to determine the most significant ones for λ . In this way, we are able to identify a set of the 10 most significant covariates, which we use in the computations to follow.

Figure 5 shows eBrAC for given TAC signals for three drinking sessions. We used a covariate-dependent transmission matrix as in (7) based on the three most significant covariates (scale-derived fitness age, hips, and waist) for λ in (14). In each case, we display eBrAC for a transmission matrix depending on the covariates of the corresponding drinking session and an eBrAC signal obtained using a transmission matrix that is independent of the covariates. Figure 5a shows only very minor differences between the covariate-dependent and -independent estimates. We note that this was the case in the vast majority of the test drinking episodes we evaluated. However, Figures 5b and 5c show examples of eBrAC where the inclusion of covariates leads to significantly different estimates.

We assume that the negligible influence of the covariates on the transitions lies in the structure of the fitted transition matrix and the model itself. If we then also assume that the hidden states represent discrete BAC levels or their surrogates, we would expect that the states would tend to either remain fixed or transition to a neighboring state or BAC level. Consequently, the trained transition matrix has tended to be either banded or close to the identity and this overall behavior was unaffected by the covariates.

In the context of the problem at hand, it seems more meaningful to include the covariates in the state-conditional emission densities as in (8). Due to the Markovian assumption, a

certain level of the unobserved BAC at a given time will be expressed as a particular TAC level. Thus, it is not unreasonable to assume that the TAC level for a given BAC level depends on covariates. Moreover, this dependence may be more readily identifiable in the data than perhaps any dependence of the transition probabilities on the covariates. Thus, in our HMM we considered covariate-dependent means for the emission distributions.

Figure 6 shows eBrAC signals for given TAC signals for three drinking sessions using covariate-dependent emission distributions as in (8) based on the three most significant covariates for λ in (14). In every case, we display an eBrAC signal for covariate-dependent emissions and an eBrAC signal for covariate-independent emissions. It can be seen that the differences between covariate-dependent and - independent estimates are larger than in the case of the transition probabilities in Figure 5, and we observed this consistently in the drinking episodes in our test set.

Since the effect of introducing covariates into the state transitions negligibly improved fit while greatly increasing the complexity of the model, in what follows we focus exclusively on covariate-dependent emissions. In particular, we are interested in quantitatively assessing the degree to which including the pre-selected set of 10 covariates improves the predictive power of our model. For every $k \in \{1, \dots, 10\}$, we train the HMM with all possible subsets of k covariates and report the log-likelihood of the fitted model together with the average eBrAC error on both the training set and the test set. For these errors, e , we only use the TAC signal of a drinking episode to obtain the eBrAC signal and then compute the squared deviation from the true BrAC signal, that is

$$e = \|\tilde{\mathcal{O}}_2 - \mathcal{O}_2\|_2 = \left(\sum_{t=1}^T (\tilde{o}_{t,2} - o_{t,2})^2 \right)^{1/2}. \quad (15)$$

Our results are in Figure 7. It is clear that including the covariates increases the log-likelihood of the fitted model. There appears to be a linear relationship between the number of covariates used and the log-likelihood. More significantly, there appears to be a similar relation when one considers the prediction errors on the training set. Including more covariates yielded a lower training error with up to 9 covariates but then at 10 covariates started to increase the training error. Apparently, increasing the number of parameters eventually leads to over-fitting that tends to outweigh the information gain from another covariate. This same behavior was also observed in the test set. Here, the error decreased with up to 8 covariates and then started to increase. Thus, it appears to be important to balance the disadvantages of increasing model complexity with the benefits to the overall quality of the fit to the data when deciding to include additional covariates.

6.5. Uncertainty Quantification

In light of the discussion in Section 4, at each time t , we use local decoding to find the probabilities $P(X_t = i | \mathcal{O}, \theta)$ for $i \in \{1, \dots, S\}$ and consider the most likely states $i_1^*, \dots, i_r^*$ such that $P(X_t \in \{i_1^*, \dots, i_r^*\} | \mathcal{O}, \theta) \geq 0.9$.

That is, to borrow a term from Bayesian statistics, we construct a conservative 90% credible interval for the decoded states. Typically, the number r of states to consider here is very small. For each of the states $k \in \{i_1^*, \dots, i_r^*\}$, we then consider the 90% credible interval $I_k = (\mu_{k2} \pm 1.645\sqrt{\Sigma_{k22}})$ for the emission. Finally, we use the maximal and minimal values of $\cup_{k=1}^r I_k$ as the upper and lower bound, respectively, for the credible band.

In Figure 8, we display the eBrAC signal based on the global decoding using the Viterbi algorithm together with the credible bands constructed as described above. The displayed drinking episodes in Figures 8a and 8b are the same as in 4d and 4f, respectively.

6.6. Results on Field Data

The BrAC and TAC for all of the drinking episodes considered thus far in this section were collected in the laboratory under controlled conditions following a prescribed dosing protocol. We now use our approach to obtain eBrAC for TAC collected under naturalistic drinking conditions in the field and seek to recover the corresponding eBrAC signal. The five naturalistic drinking episodes were collected by one of the authors (S.E.L.), with the first four episodes using the WrisTASTM7 transdermal alcohol biosensor (Giner, Inc., Waltham, Massachusetts) and the fifth episode using the SCRAM-CAM. The HMM was trained with the SCRAM-CAM data set discussed earlier.

Figure 9 shows the results from these naturalistic drinking sessions. It can be seen that the eBrAC signal captures most features of the true BrAC signal quite accurately. Although the peak eBrAC level in some cases differs from the true peak BrAC level, the shape of the eBrAC signal nicely matches the true BrAC signal. In particular, the fact that the peak of the BrAC signal appears earlier than the peak of the TAC signal is well captured.

Our data-intensive approach only worked well on field data if the natural drinking episodes, to some degree, resembled the drinking episodes on which the HMM was trained. Not included in Figure 9 are the naturalistic drinking sessions we tested our approach on with very low peak TACs. These single drink episodes had peak TACs of around .020 which is much lower than the peak TACs of our laboratory collected training episodes that were dosed to reach a peak BrAC of .080. Since there were no drinking episodes with a very low peak TAC present in the training data, the model could not learn what BrAC values are associated with low peak TAC signals and consequently produced false eBrACs for these sessions. In Figure 9(e) we see that our HMM also failed to capture the bi-modal nature of this particular drinking episode. We again believe that this may be attributed to the lack of multi-modal drinking sessions in the data available to us for the proof-of-concept study being reported on here. A training set with a broader range of naturalistic drinking sessions, including those with lower, higher, and multi-modal levels of consumption would test the fit of our HMM under more varied drinking conditions. This is an important direction for future research.

6.7. Results for the Physics-Informed Viterbi Algorithm

So far, all the results shown were obtained using the classical Viterbi algorithm for the BrAC estimation. Now, we want to compare this approach to the physics-informed Viterbi

algorithm outlined in Section 5. Figure 10 shows the BrAC signals for three drinking episodes from the test set together with eBrAC signals that were obtained using the classical and the physics-informed Viterbi algorithm. In Figure 10a, both eBrAC signals align for most of the drinking episode. However, the eBrAC signal that was obtained without the physics regularization shows an artifact towards the end of the drinking episode. This artifact is not physically meaningful as the drinking episode is not bi-modal in nature. Interestingly, the physics-informed eBrAC signal does not exhibit this artifact; here the physical regularization effectively helped to produce a physically meaningful estimate. A similar situation can be seen in Figure 10b. The eBrAC signal without the physics regularization is unrealistically long and has an artifact towards the end. The physics-informed signal better captures the true episode, although in this case the estimated signal is too short. A more extreme case is shown in Figure 10c: without the physics information, the eBrAC signal hardly shows any drinking. Only the physics-informed eBrAC signal captures length and intensity of the drinking episode approximately correct. Although we only display three episodes from the test set, we observe that the physical regularization effectively helps to avoid these artifacts when they appear in the estimates signals.

6.8. Correlation Between Hidden States and BAC

As indicated in Section 2, one could ask whether there is sufficient evidence to identify a correlation between the hidden states and, for example, BAC. Based on the assumption that BAC would be more highly correlated with BrAC than with TAC, we ordered the hidden states based on the mean of the BrAC component of the emission distributions. In Figure 11, we display the TAC and BrAC signal of two drinking sessions together with the sequence of ordered hidden states based on the TAC signal and the Viterbi algorithm. From these plots, we observe that early in the episode the hidden states quickly reach their peak at essentially the same time the BrAC signal reaches its peak. In the later part of the episode, however, the hidden states tend to decay more slowly with the TAC following a parallel but somewhat delayed trajectory. Finally, when the hidden states have reached zero, the TAC follows suit and decays to zero as well. To conclude from the limited evidence provided here that the hidden states are, in fact, BAC, one would have to argue that when BAC drops below approximately .03 percent alcohol, the breath analyzer is no longer able to detect the presence of alcohol, which is a conclusion that would be at odds with experience. It seems more likely that the hidden states are simultaneously capturing a number of the mechanisms the body uses to metabolize and excrete alcohol, or at least those that can be observed in breath and transdermal data. More research is required to better understand the hidden states and their relationship among BAC, BrAC, TAC, and covariates.

7. Summary and Concluding Remarks

In this paper we have developed a novel method for obtaining an estimate of blood or breath alcohol concentration from biosensor measured transdermal alcohol level. Our scheme is based on first training a two emission (BrAC and TAC) variable hidden Markov model using available contemporaneous BrAC and TAC signals from a population cohort and the Baum-Welch algorithm. We then supply the trained model with a given subject's TAC measurements only and use the Viterbi algorithm to determine the most likely path through

the hidden Markov chain. Finally, the estimated BrAC derived from the subject's observed TAC is taken to be the mean BrAC emissions corresponding to the Viterbi path through the chain. In addition we have allowed for the dependence of the model parameters (i.e. the Markov chain transition and initial probabilities and the emission variable distribution parameters) on relatively readily observable subject-level and episode-level covariates in the form of (multinomial logistic for transition and initial probabilities, standard linear for emission distributions) regression models that are trained at the same time as the underlying HMM. We developed a method based on local state decoding that allows for the quantification of the uncertainty in the estimated BrAC, and we showed how a decoding of the hidden states can be used to gain insight into the body's internal processing of ethanol in the form of estimated BAC. Finally, we have proposed an innovative physics-informed Viterbi algorithm that takes into account both the stochastic model and the prior information in the form of a PDE model. Our numerical results show that the physics-informed HMM presented herein is able to yield meaningful BrAC estimates and that the corresponding uncertainty can be quantified by leveraging the stochastic nature of the model. It was also shown that the inclusion of covariates leads to lower training and test errors, and that the physics-informed Viterbi algorithm helps to produce physically meaningful BrAC estimates without unrealistic artifacts.

Moving forward, there are two areas of further research that we are pursuing. The first focuses on the application of the HMM on a larger and more diverse dataset. One particular strength of this HMM approach is that the model can be trained on very large datasets. However, so far the available dataset only contains 256 drinking episodes from a laboratory controlled setting. It would be interesting to see how well the proposed method works on a larger and more varied dataset that contains multi-modal drinking episodes and more naturalistic drinking settings.

A second area of further research deals with the physics-based regularization described in Section 5. Although we have a deterministic physics-based model that has been shown in our earlier work [10, 15, 43] to be able to capture the dynamics of the transdermal transport of ethanol, the parameters in the model tend to depend on the individual wearing the sensor and ambient environmental conditions. Our approach incorporates this individual model into a population model. The question is how the values of the parameters in the PDE model should be chosen. In our proof-of-concept study here, we simply set parameters equal to their population means. However, we have developed more sophisticated PDE-based population models where the parameters are assumed to be random with their distributions fit to their values over the entire population. We have then shown how these random PDEs can be transformed into an equivalent weak form wherein the random parameters are now treated as additional spatial variables (see, for example [48, 49, 50, 58]). We intend to investigate how these more sophisticated population models perform when introduced into our HMM model.

In addition, while the physics-based PDE model we used here captures only the transdermal transport of ethanol, the states in the hidden Markov chain in the HMM are intended to capture much more of the body's processing of alcohol including, for example, ingestion, absorption through the gut, metabolism by the liver, etc. We have also developed first

principles physics-based semi-linear hybrid PDE/ODE models to capture these dynamics as well, including the nonlinear Michaelis-Menten enzyme catalyzed kinetics of the metabolism of alcohol in the liver. We are also planning to investigate how a physics-informed Viterbi algorithm using this model with random parameters affects accuracy.

Acknowledgments

This research supported in part by the National Institute on Alcohol Abuse and Alcoholism (NIAAA) grant no. R01AA026368.

Supported in part by a grant from USC's Women in Science and Engineering (WiSE) Undergraduate Research Experience Program.

References

- [1]. Banks HT, Flores K, Rosen I, Rutter E, Sirlanci M, and Thompson W, "The Prohorov metric framework and aggregate data inverse problems for random PDEs", *Commun. Appl. Anal*, 22, 415–446, 2018. [PubMed: 35958041]
- [2]. Baum LE and Petrie T, "Statistical Inference for Probabilistic Functions of Finite State Markov Chains", *Ann. Math. Statist* 37(6), 1554–1563, 1966.
- [3]. Baum LE and Eagon JA, "An inequality with applications to statistical estimation for probabilistic functions of Markov processes and to a model for ecology", *Bull. Amer. Math. Soc* 73(3), 360–363, 1967.
- [4]. Baum LE and Sell GR, "Growth transformations for functions on manifolds", *Pacific J. Math* 27(2), 211–227, 1968.
- [5]. Baum LE, Petrie T, Soules G, and Weiss N, "A Maximization Technique Occurring in the Statistical Analysis of Probabilistic Functions of Markov Chains", *Ann. Math. Statist* 41(1), 164–171, 1970.
- [6]. Baum LE, "An inequality and associated maximization technique in statistical estimation for probabilistic functions of Markov processes", *Inequalities*, vol. 3, 1–8, 1972.
- [7]. Cappe O, "Ten years of HMMs", URL: <https://www.ai.rug.nl/vakinformatie/sr/articles/hmmbibliography-1989-2000>, 2001.
- [8]. Cappe O, Moulines E, and Ryden T, *Inference in Hidden Markov Models*, Springer, New York, 2005.
- [9]. Casella G and Berger RL, *Statistical Inference*, Duxbury, Pacific Grove, CA, 2002.
- [10]. Dai Z, Rosen IG, Wang C, Barnett N, and Luczak SE, "Using drinking data and pharmacokinetic modeling to calibrate transport model and blind deconvolution based data analysis software for transdermal alcohol biosensors", *Math. Biosci. Eng.*, 13, 911–34, 2016. [PubMed: 27775390]
- [11]. Dempster AP, Laird NM, and Rubin DB, "Maximum likelihood from incomplete data via the EM algorithm", *J. of the Royal Stat. Soc., Series B* 39, 1–38, 1977.
- [12]. Dougherty DM, Charles NE, Acheson A, John S, Furr RM, and Hill-Kapturczak N, "Comparing the detection of transdermal and breath alcohol concentrations during periods of alcohol consumption ranging from moderate drinking to binge drinking", *Exp. Clin. Psychopharm*, 20, 373–81, 2012.
- [13]. Dougherty DM, Karns TE, Mullen J, Liang Y, Lake SL, Roache JD, and Hill-Kapturczak N, "Transdermal alcohol concentration data collected during a contingency management program to reduce at-risk drinking", *Drug and Alc. Dep.*, 148, pp. 77–84, 2015.
- [14]. Drake AW, "Discrete-state Markov processes", Chapter 5 in *Fundamentals of Applied Probability Theory*, New York, McGraw-Hill, 1967.
- [15]. Dumett MA, Rosen IG, Sabat J, Shaman A, Tempelman L, Wang C, and Swift R, "Deconvolving an estimate of breath measured blood alcohol concentration from biosensor collected transdermal ethanol data", *Appl. Math. Comp.*, 196, 724–743, 2008.

- [16]. Fairbairn C, Kang D, and Bosch N, "Using machine learning for real-time BAC estimation from a new-generation transdermal biosensor in the laboratory", *Drug and Alc. Dep*, 216, 108–205, 2021.
- [17]. Feller W, *An Introduction to Probability Theory and Its Applications I*, Wiley, New York, 1968.
- [18]. Hawekotte K, Luczak SE, and Rosen IG, "Deconvolving breath alcohol concentration from biosensor measured transdermal alcohol level under uncertainty: A Bayesian approach", *Math. Biosci. Eng.*, to appear, 2021.
- [19]. Hill-Kapturczak N, Roache JD, Liang Y, Karns TE, Cates SE, and Dougherty DM, "Accounting for sex-related differences in the estimation of breath alcohol concentrations using transdermal alcohol monitoring", *Psychopharmacology*, 232, 115–23, 2015.
- [20]. Hill-Kapturczak N, Lake SL, Roache JD, Cates SE, Liang Y, and Dougherty DM, "Do variable rates of alcohol drinking alter the ability to use transdermal alcohol monitors to estimate peak breath alcohol and total number of drinks?", *Alc. Clin. Exp. Res.*, 38, 2517–2522, 2014.
- [21]. Jones AW, "Determination of liquid/air partition coefficients for dilute solutions of ethanol in water, whole blood, and plasma", *J. Anal. Toxicol.*, 7, 193–197, 1983. [PubMed: 6101261]
- [22]. Jones A and Andersson L, "Comparison of ethanol concentrations in venous blood and end-expired breath during a controlled drinking study", *Forensic Sci. Int.*, 132, 18–25, 2003. [PubMed: 12689747]
- [23]. Juang BH and Rabiner LR, "Hidden Markov Models for Speech Recognition", *Technometrics*, vol. 33, no. 3, pp. 251–72, 1991.
- [24]. Karniadakis G, Kevrekidis I, Lu L, Perdikaris P, Wang S and Yang L, "Physics-informed machine learning", *Nat. Rev. Phys.*, 3, 422–440, 2021.
- [25]. Labianca DA, "The chemical basis of the breathalyzer: A critical analysis", *J. Chem. Educ.*, 67, 259–61, 1990.
- [26]. Liu T, Lemeire J, and Yang L, "Proper initialization of hidden Markov models for industrial applications", 2014 IEEE China Summit and International Conference on Signal and Information Processing (ChinaSIP), Xi'an, China, 490–494, 2014.
- [27]. Liu W, Lai Z, Bacsa K, and Chatzi E, "Physics-guided Deep Markov Models for Learning Nonlinear Dynamical Systems with Uncertainty", arXiv:2110.08607v1 [cs.LG], 16 Oct 2021.
- [28]. Lourenc O, Mattila R, Rojas C, Hu X, and Wahlberg B, "Hidden Markov models: inverse filtering, belief estimation and privacy protection", *J. Syst. Sci. Complex* 34, 1801–1820, 2021.
- [29]. Luczak SE, and Ramchandani VA. "Special issue on alcohol biosensors: Development, use, and state of the field: summary, conclusions, and future directions", *Alc.*, 81, 161–165, 2019.
- [30]. Mattila R, Rojas C, Krishnamurthy V, and Wahlberg B, "Inverse filtering for hidden markov models With applications to counter-adversarial autonomous systems", *IEEE Trans. Signal Process.*, 68, 4987–5002, 2020.
- [31]. MacKay D, "Chapter 20. An Example Inference Task: Clustering", *Information Theory, Inference and Learning Algorithms*, Cambridge University Press, 284–292, 2003.
- [32]. Moore S, Radunskaya A, Zollinger E, Grant K, Gonzales S, and Baker E, "Time for a drink? a mathematical model of non-human primate alcohol consumption", *Front. Appl. Math. Stat.*, 22, 2019.
- [33]. Mor B, Garhwal S, and Kumar A, "A systematic review of hidden Markov models and their applications", *Arch. Comput. Methods Eng* 28, 1429–1448, 2021.
- [34]. Murphy KP, "Fitting a conditional linear Gaussian distribution", 2000. Available at: <https://www.researchgate.net/publication/2453789>.
- [35]. Oszkinat C, Luczak SE, and Rosen IG, "Uncertainty quantification in the estimation of blood alcohol concentration using physics-informed neural networks", *IEEE Trans. Neural Netw. Learn. Syst.*, to appear, 2022.
- [36]. Oszkinat C, Luczak SE, and Rosen IG, "A physics-informed long short-term memory network based on an abstract parabolic system for the estimation of breath alcohol concentration from transdermal alcohol biosensor data", in preparation, 2022.
- [37]. Pazy A, *Semigroups of Linear Operators and Applications to Partial Differential Equations*, New York, Springer, 1983

- [38]. Rabiner LR, “A tutorial on hidden Markov models and selected applications in speech recognition”, *Proc. IEEE*, 77, 257–286, 1989.
- [39]. Rabiner LR, Levinson SE, and Sondhi MM, “On the application of vector quantization and hidden Markov models to speaker-independent, isolated word recognition”, in *The Bell System Technical Journal*, vol. 62, no. 4, pp. 1075–1105, 1983.
- [40]. Raissi M, “Deep hidden physics models: deep learning of nonlinear partial differential equations”, *J. Mach. Learn. Res*, 19, 932–955 2018.
- [41]. Raissi M and Karniadakis G, “Hidden physics models: machine learning of nonlinear partial differential equations”, *J. Comp. Phys*, 357, 2017.
- [42]. Raissi M, Perdikaris P, and Karniadakis G, “Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations”, *J. Comp. Phys*, 378, 2018.
- [43]. Rosen IG, Luczak SE, and Weiss J, “Blind deconvolution for distributed parameter systems with unbounded input and output and determining blood alcohol concentration from transdermal biosensor data”, *Appl. Math. Comp*, 231, 357–376, 2014.
- [44]. Saldich EB, Wang C, Rosen IG, Goldstein L, Bartroff J, Swift RM, and Luczak SE, “Obtaining high-resolution multi-biosensor data for modeling transdermal alcohol concentration data”, *Alc.: Clin. Exp. Res*, 44, 181A, 2020.
- [45]. Schwartz R, Chow Y, Roucos S, Krasner M, and Makhoul J, “Improved hidden Markov modeling of phonemes for continuous speech recognition”, *ICASSP '84. IEEE International Conference on Acoustics, Speech, and Signal Processing*, pp. 21–24, 1984.
- [46]. Shin Y, Darbon J, and Karniadakis GE, “On the convergence of physics informed neural networks for linear second-order elliptic and parabolic type PDEs”, *arXiv:2004.01806 [math.NA]*, 2020.
- [47]. Shirley KE, Small DS, Lynch KG, Maisto SA, and Oslin DW, “Hidden Markov models for alcoholism treatment trial data”, *Ann. Appl. Stat*, 4, 366–395, 2010.
- [48]. Sirlanci M, Luczak SE, and Rosen IG, “Approximation and convergence in the estimation of random parameters in linear holomorphic semigroups generated by regularly dissipative operators”, *Proc. 2017 ACC*, 3171–3176, 2017.
- [49]. Sirlanci M, Luczak SE, and Rosen IG, “Estimation of the distribution of random parameters in discrete time abstract parabolic systems with unbounded input and output: Approximation and convergence”, *Comm. Appl. Anal*, 23, 287–329, 2019.
- [50]. Sirlanci M, Rosen IG, Luczak SE, Fairbairn CE, Bresin K, and Kang D, “Deconvolving the input to random abstract parabolic systems; A population model-based approach to estimating blood/breath alcohol concentration from transdermal alcohol biosensor data”, *Inv. Prob*, 34, 125006, 2018.
- [51]. Sirlanci M, Luczak SE, Fairbairn CE, Kang D, Pan R, Yu X, and Rosen IG, “Estimating the distribution of random parameters in a diffusion equation forward model for a transdermal alcohol biosensor”, *Automatica*, 106, 101–109, 2019. [PubMed: 31814628]
- [52]. Swift RM, “Transdermal alcohol measurement for estimation of blood alcohol concentration”, *Alc.: Clin. Exp. Res*, 24, 422–423, 2000.
- [53]. Tanabe H, *Equations of Evolution*. San Francisco: Pitman, 1979.
- [54]. Viterbi A, “Error bounds for convolutional codes and an asymptotically optimum decoding algorithm”, *IEEE Trans. Inf. Th*, 13, 260–269, 1967.
- [55]. Witkiewitz K, Maisto S, and Donovan D, “A comparison of methods for estimating change in drinking following alcohol treatment”, *Alc. Clin. Exp. Res*, 34, 2116–2125, 2010.
- [56]. Yang Y and Perdikaris P, “Adversarial uncertainty quantification in physics-informed neural networks”, *J. Comp. Phys*, 394, 136–152, 2019.
- [57]. Zucchini W, MacDonald IL, and Langrock R, *Hidden Markov Models for Time Series: An Introduction Using R*, 2nd ed., CRC Press, 2016.
- [58]. Yao M, Luczak S, Saldich E and Rosen I, “A Population Model-Based LQG Compensator for the Control of Intravenously-Infused Alcohol Studies and Withdrawal Symptom Prophylaxis Using Transdermal Sensing”, submitted, 2022.



Figure 1:
(Left) Giner, Inc. WristAS™7 transdermal continuous alcohol monitoring device, (Right)
SCRAM Systems transdermal continuous alcohol monitoring device.




Figure 2: (Upper) Typical covariate-informed, two-phase, training (left), and prediction/inversion/deconvolution (right) system for obtaining breath alcohol concentration from transdermal alcohol concentration; (Lower) Our covariate-dependent, physics-informed, HMM-based, two-phase, training (left), and prediction/inversion (right) system for obtaining breath alcohol concentration from transdermal alcohol concentration.




Figure 3:
Basic architecture of an HMM.




Figure 4:

Performance of the HMM on three drinking episodes from the test set. The figures in the upper row display TAC and BrAC data together with their predictions when both TAC and BrAC are used as input for the Viterbi algorithm. In the lower row, only TAC is used as an input to the Viterbi algorithm to produce the eBrAC signal.




Figure 5: eBrAC signals based on the TAC signal and covariate-dependent transition probabilities for test set episodes. The figures show BrAC data, eBrAC using a covariate-dependent transition matrix, and eBrAC using a covariate-independent transition matrix.




Figure 6: eBrAC signals based on the TAC signal in the case of covariate-dependent linear Gaussian emissions for three drinking episodes from the test set. The figures display BrAC data, eBrAC using a covariate-dependent conditional Gaussian emission, and eBrAC using a covariate-independent conditional Gaussian emission.




Figure 7: Log-likelihood of the trained model together with BrAC estimation error (15) for the training and test sets vs the number of included covariates. For any number k of covariates, all possible combinations of k covariates out of 10 have been considered and averaged.




Figure 8:

Demonstration of the proposed uncertainty quantification for the BrAC estimation for three drinking episodes from the test set. The figures display BrAC data, eBrAC based on the global state decoding using the Viterbi algorithm, and a credible band for the BrAC signal based on the uncertainty quantification.




Figure 9:
Application of the fit HMM model to natural drinking episodes recorded in the field. The figures show recorded TAC and BrAC signals and eBrAC based only on the TAC signal.




Figure 10:

Performance of the physics-informed Viterbi algorithm on three drinking episodes from the test set. The figures depict the true BrAC data as well as eBrAC signals that are based on the TAC only. The black lines show eBrAC obtained from the physics-informed Viterbi algorithm, whereas the red lines show the estimates obtained using the classical Viterbi algorithm.




Figure 11:
TAC and BrAC together with the decoded hidden states based on the TAC.

Table 1:

Median (computed with respect to all training (TR) and testing (T) data sets) relative error in eBrAC and eTAC signals, peak eBrAC and eTAC values, time of peak eBrAC and eTAC, and area underneath eBrAC and eTAC curves with eBrAC and eTAC estimated using either BrAC and TAC or on TAC alone.

	eBrAC from BrAC and TAC				eBrAC from TAC Only			
	rL_2	rP	$rAUC$	$ \Delta T $	rL_2	rP	$rAUC$	$ \Delta T $
TAC _{Tr}	.2133	.1427	.0594	.3200	.1381	.1039	.0330	.1900
TAC _T	.2127	.1207	.0674	.3850	.1419	.1220	.0350	.3700
BrAC _{Tr}	.2215	.0717	.0563	.2850	.4975	.2320	.2659	.3000
BrAC _T	.2447	.0842	.0735	.2650	.4855	.2175	.2529	.2800