Learning Outcome-Discriminative Dynamics in Multivariate Physiological Cohort Time Series

Shamim Nemati[1], Li-wei H. Lehman[2] and Ryan P. Adams[1]

Abstract— Model identification for physiological systems is complicated by changes between operating regimes and measurement artifacts. We present a solution to these problems by assuming that a cohort of physiological time series is generated by switching among a finite collection of physiologically-constrained dynamical models and artifactual segments. We model the resulting time series using the switching linear dynamical systems (SLDS) framework, and present a novel learning algorithm for the class of SLDS, with the objective of identifying time series dynamics that are predictive of physiological regimes or outcomes of interest. We present exploratory results based on a simulation study and a physiological classification example of decoding postural changes from heart rate and blood pressure. We demonstrate a significant improvement in classification over methods based on feature learning via expectation maximization. The proposed learning algorithm is general, and can be extended to other applications involving state-space formulations.

I. INTRODUCTION

Physiological control systems involve multiple interacting variables operating in feedback loops that enhance an organism's ability to self-regulate and respond to internal and external disturbances. The resulting multivariate time series often exhibit rich dynamical patterns that are altered under pathological conditions, and are therefore informative of health and disease [1], [2], [3], [4]. Using nonlinear [1], [2] indices of heart rate (HR) variability (i.e., beat-to-beat fluctuations in HR), researchers have shown that subtle changes to the dynamics of HR may act as an early sign of adverse cardiovascular outcomes (e.g., mortality after myocardial infarction [3]) in large cohort studies. However, these studies fall short of assessing the multivariate dynamics of the vital signs (e.g., heart rate, blood pressure, respiration, etc.), and do not yield any mechanistic hypotheses for the observed deteriorations of normal variability. This shortcoming is in part due to the inherent difficulty of parameter estimation in physiological time series, where one is confronted by nonlinearities (including rapid regime changes), measurement artifacts, and/or missing data, which are particularly prominent in ambulatory recordings (due to patient movements) and bedside monitoring (due to equipment malfunction).

In previous work [5], [6], we developed a framework for automated discovery of shared dynamics in multivariate physiological time series from large patient cohorts. A central premise of our approach was that even within heterogeneous cohorts (with respect to demographics, genetic factors, etc.) there are common "phenotypic" dynamics that a patient's vital signs may exhibit, reflecting underlying pathologies (e.g., detraction of the baroreflex system) or temporary physiological state changes (e.g., postural changes or sleep/wake related changes in physiology). We used the switching linear dynamical system (SLDS) framework to automatically segment the time series into regions with similar *dynamics*, i.e., time-dependent rules describing the evolution of the system state. Importantly, the framework allows for incorporation of physiologically-constrained linear models (e.g., via linearization of the nonlinear dynamics around equilibrium points of interest) to derive mechanistic explanations of the observed dynamical patterns, for instance, in terms of directional influences among the interacting variables (e.g., baroreflex gain or chemoreflex sensitivity).

Although we assumed *a priori* knowledge of the underlying physiology to constrain the dynamical models, the model parameters have to be learned from the data. As noted earlier, artifacts in physiological recordings and incomplete knowledge of the underlying physiology may hinder system identification using traditional approaches such as maximum likelihood estimation. In this work, we propose a learning algorithm specifically designed to learn dynamical features of data that are predictive of patient outcomes such as a patient's physiological state or long-term survival. Previous approaches have used a two-stage procedure: unsupervised feature extraction followed by supervised learning for outcome discrimination. Here we take the novel approach of *jointly* learning the dynamics and the classifier. Rather than depending on label-free unsupervised learning to discover relevant features of the time series, we build a system that expressly learns the dynamics that are most relevant for predicting patient outcome.

II. METHODS

Assume we are given a collection of N multivariate time series and the associated outcome variables: $\{(y^{(1)}, O^{(1)}), (y^{(2)}, O^{(2)}), \cdots, (y^{(N)}, O^{(N)})\}$, where the n-th time series $y^{(n)}$ is of length T_n, and may include M channels. The corresponding label $O^{(n)}$ can be a scalar such as a

(a) Simulated time series
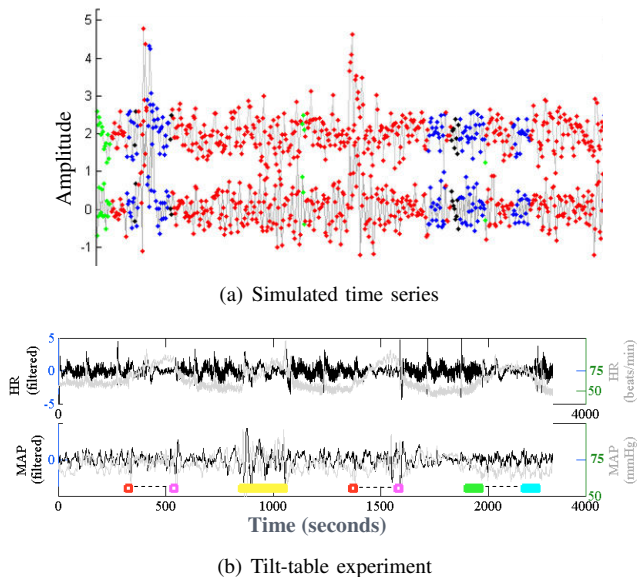

(b) Tilt-table experiment

Fig. 1. **(a)** An example bivariate time series out of the 200 simulated. The time series were divided into 4 categories, each statistically having different proportions of four modes ($M_1, \cdots, M_4$, color-coded as red, blue, green, and black, respectively). Here, we introduced an offset of 2 in one of the channels of each time series to improve visibility. **(b)** An example of heart rate and mean blood pressure from the tilt-table experiment is shown in the bottom panel (actual values in gray and filtered values in black). The ground truth segmentation is indicated using color-coded horizontal lines – green to cyan: slow tilt up and down to supine; red to pink: rapid tilt up and down to supine; yellow: standing up and back to supine.

discrete patient outcome, or it may itself be a length-T_n time series vector that assigns a label to each instant. Our objective is to find shared dynamical features across the different time series that are predictive of the labels.

A. Datasets

1) Simulated time series with Switching Dynamics: We simulated 200 bivariate time series with dynamic switching among four modes ($J = 4$, color-coded in Fig. 1(a)). All four dynamical modes were stable bivariate ($M = 2$) autoregression (AR) models of order two ($P = 2$). To increase the heterogeneity of the dataset, the time series were simulated using four different Markov transition matrices (the stationary distribution of the four categories were [0.67, 0.10, 0.10, 0.13], [0.14, 0.57, 0.19, 0.10], [0.08, 0.16, 0.54, 0.22], and [0.09, 0.09, 0.23, 0.59]). Additionally, we introduced approximately 10% variation in the AR coefficients across each realization by adding white Gaussian noise with standard deviation 0.05 to each of the AR coefficients. Finally, all time series included two randomly-placed large-amplitude artifacts (uniform random noise in the interval of [0,15]) of 10 samples duration. Fig. 1(a) shows an example of the simulated time series.

2) Tilt-Table Experiment: Time series of HR and mean arterial blood pressure (MAP) were acquired from 10 healthy subjects undergoing a tilt-table experiment. The details of the protocol are described in Heldt et al. [7]. Briefly, subjects were placed in the supine position and secured to a table. Tilting was performed at various speeds from the horizontal position to the vertical position and back to supine, generating four postural categories of (1) supine, (2) slow-tilt,

(3) fast tilt, and (4) standing. One example of the resulting time series is shown in Fig. 1(b).

Data Pre-processing: Since we were interested in the interaction between HR and MAP in the frequency range pertinent to sympathetic and parasympathetic regulation [5], the time series of HR and MAP were high-pass filtered to remove the steady-state baseline and any oscillation in the time series with a period slower than 100 beats. This filtering was done using a 7th order Butterworth digital filter with cutoff frequency of 0.01 cycles/beat.

B. Learning Switching Dynamics in Cohort time series

Switching Linear Dynamical Systems: The switching linear dynamical system (SLDS) framework [8] models time series using two layers of evolution. In the high-level layer, the time series evolves through a set of J modes according to Markovian dynamics. In the lower level, each of these modes corresponds to a unique linear dynamical system that evolves a continuous state and produces the observed time series. The generative model is as follows: a latent process for each time series $S_t^{(n)} \in \{1, \cdots, J\}$ evolves according to Markovian dynamics with initial distribution $\pi^{(n)}$ and $J \times J$ transition matrix Z. Each of the n series has an unobserved continuous state variable $x_t^{(n)} \in \mathbb{R}^M$ that evolves according to linear dynamics which are determined by the current mode $S_t^{(n)}$, and produces observations $y_t^{(n)}$. The jth linear system has state dynamics $A^{(j)}$, observation matrix $C^{(j)}$, state noise covariance $Q^{(j)}$, and observation noise covariance $R^{(j)}$:

$$x_t^{(n)} = A^{(S_t^{(n)})} x_{t-1}^{(n)} + v_t \qquad v_t \sim N(0, Q^{(S_t^{(n)})})$$
$$y_t^{(n)} = C^{(S_t^{(n)})} x_t^{(n)} + w_t \qquad w_t \sim N(0, R^{(S_t^{(n)})}).$$

We refer to these mode-specific dynamics together as $\Theta^{(j)}$.

EM for Parameter Learning in Switching Dynamical Systems: A comprehensive treatment of the EM algorithm for SLDS is presented in Murphy (1998) [8]. Briefly, in practice we neither know the set of switching variables nor the parameters that define the modes. EM is a two-pass iterative
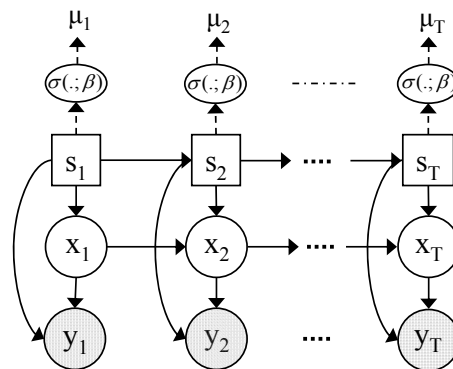

Fig. 2. time series dynamics-based sequential labeling. The first three layers (from bottom) depict a graphical model representation of the switching linear dynamical system, and the last layer estimates parameters of multinomial probability of outcomes ($\mu_1, \cdots, \mu_T$); via the multinomial logistic regression function $\sigma(.;\beta)$, with parameters β. The round nodes are continuous and Gaussian random variables, the square nodes are discrete random variables, and the elliptical nodes are deterministic functions. Shaded nodes are observed and the rest are hidden. Solid arrows denote the conditional dependencies among the random variables.
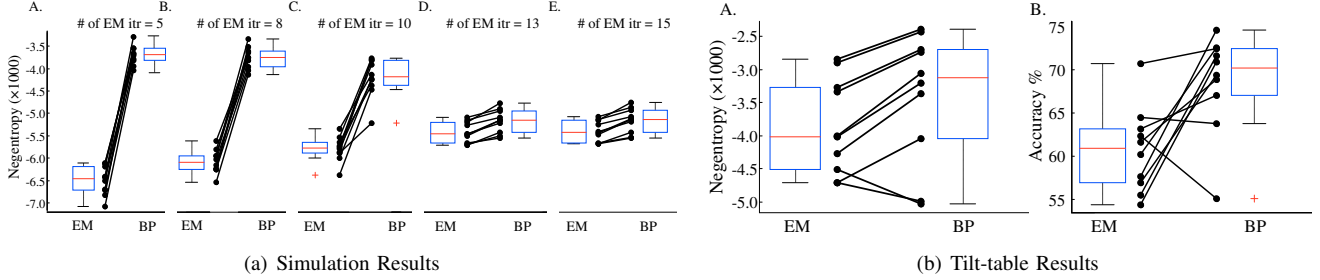
(a) Simulation Results

(b) Tilt-table Results

Fig. 3. **(a)** Panels A-E show classification performance over ten folds using the EM for 5, 8, 10, 13, and 15 iterations, and EM followed by 30 iterations of supervised learning (BP). **(b)** Panel A shows 10-fold cross-validation results, comparing the accuracy achieved using 10 iteration of EM versus 10 iterations of EM followed by 30 iterations of the L-BFGS optimizer. Panel B shows a comparison of the two techniques in terms of classification accuracy.

algorithm: (1) in the expectation (E) step we obtain the expected values of the latent variables $\{\{x_t^{(n)}, S_t^{(n)}\}_{t=1}^{T_n}\}_{n=1}^{N}$ using a modified Kalman smoother [8], and (2) in the maximization (M) step we find the model parameters $\{\Theta^{(j)}\}_{j=1}^{J}$, Markov dynamics Z and initial conditions $\pi^{(n)}$ that maximize the expected complete data log likelihood. In our implementation of the EM algorithm, we achieve *shared* dynamics by pooling together all subjects' inferred variables in the M step. We impose physiological constraints on the model parameters using a constrained least square approach. Iteration through several steps of the EM algorithm results in learning a set of J shared modes and a global transition matrix Z for all the patients.

Sequential Labeling: After using the EM algorithm to estimate model parameters, the E-step estimates of switching variables $S_t^{(n)}$ can be used to predict the training labels $O_t^{(n)}$. Here we assume that each label $O_t^{(n)}$ can take on any one of K possible outcomes, and can be modeled using a softmax classifier with parameters β. We take the classification objective to be the negative log likelihood (*negentropy*) of the outcome labels, given the time series:

$$-\log \Pr(\mathbf{O}|\mu(\Theta,Z,\beta)) = -\sum_{n=1}^{N}\sum_{t=1}^{T_n}\sum_{k=1}^{K} \mathbf{O}_{t,k}^{(n)} \log \mu_{t,k}^{(n)}(\Theta,Z,\beta).$$
(1)

where $\mathbf{O}$ denotes the set of all the outcomes, and the estimated multinomial outcome probabilities, as a function of the various SLDS parameters, are denoted $\mu(\Theta,Z,\beta)$ (see Fig. 2).

Outcome-Discriminative Learning: Within the EM framework, unsupervised learning of the dynamics is treated separately from the discriminative learning of a mapping between switching states and outcome labels. Our objective is to design purely-supervised learning algorithm that discovers dynamical features in series that are predictive of the outcome variables. The key insight of the proposed learning algorithm is that the gradient of the objective calculated in Eq. (1) can be backpropagated through the network architecture depicted in Fig. 2 to efficiently calculate the gradient with respect to model parameters. The details of the *analytical* expressions for the involved derivatives, which allows a two-pass algorithm for calculating the exact gradients, are given in Nemati [9].

The objective of the proposed leaning algorithm is therefore to minimize the cost function in Eq. (1):

$$\Theta^\star, Z^\star, \beta^\star = \underset{\Theta, Z, \beta}{\arg\min} \{-\log \Pr(\mathbf{O}|\mu(\Theta, Z, \beta))\} \quad (2)$$

subject to the constraints that covariance matrices $Q^{(j)}$ and $R^{(j)}$ remain positive definite, and all the elements of Z stay nonnegative and each row sums to one.

1) EM-based Initialization: Due to the large number of free parameters in Eq. (2), it is necessary to carefully manage the optimization procedure in order to avoid overfitting and local minima. We initialize the parameters by running a few iterations of the EM, and then switch to a nonlinear optimization procedure, such as the L-BFGS quasi-Newton method.

III. RESULTS

A. Simulation Study

We assumed that the number of modes and the model order is known *a priori*, and tested the performance of both the EM and the proposed learning algorithm on the classification problem of labeling each time series sample as belonging to one of four modes. The results summarized in Fig. 3(b) show that the classification performance using EM improves up to 13 iterations, but further EM iterations do not improve performance on the held-out test data. Notably, the figure demonstrates the dependence of the proposed joint supervised learning on the EM initialization. In particular, supervised learning benefits from initialization with up to 8 iterations of the EM; further EM iterations seem to lower its performance, presumably due to local minima and overfitting of artifacts.

B. Tilt-Table Experiment

We constructed a sequential labeling/classification task involving the four maneuvers depicted in Fig. 1(b). We used four modes, each corresponding to an AR model of order three, to model the bivariate time series of heart rate and blood pressure. The supervised learning algorithm was initialized using 10 iterations of the EM algorithm, followed by 30 iterations of L-BFGS. The results shown in Fig. 3(b) indicate that the joint supervised learning slightly improves the negentropy cost function. We also considered the classification accuracy, which shows a significant improvement (Wilcoxon Signed-Rank Test, $p < 0.05$) as a consequence of joint supervised learning.

1) Physiological Interpretation of the Discovered Dynamics: Since we modeled the dynamics using multivariate AR models, we were able to derive the parametric power spectra corresponding to the individual channels of each time series [4]. We observed a progressive increase in the ratio of the low frequency (LF: periods of 6-20 beats) to the high frequency (HF: periods of 2-5 beats) power of the HR time series (also know as the LF/HF ratio; an index of sympathovagal activation) from supine to slow tilting, fast tilting, and standing. This indicates increased sympathetic modulations. These results were obtained by (1) calculating the parametric power spectrum of the HR for each mode, using its AR coefficients, and (2) calculating a weighted average of the HR spectrum within the segments corresponding to each postural regime, where the weights were given by the probabilities of belonging to a given mode. The estimated increase in LF/HF ratio from supine to standing was significant with both learning techniques (EM: 4.6 [4.3, 5.4] to 8.4 [8.3, 8.6] † , supervised: 4.4 [3.7 4.8] to 5.53 [5.2 6.5] †, median [interquartiles]; $\dagger$ indicates $p < 0.05$ using Kruskal-Wallis nonparametric ANOVA test).

IV. DISCUSSION AND FUTURE DIRECTION

We presented a novel technique for discriminative learning of dynamics in cohort time series. The main idea of our approach was to provide the learning algorithm with the outcomes (labels) corresponding to each time series sample (e.g., supine, slow-tilt, etc), and to learn switching linear dynamics that are maximally discriminative. Using simulated time series, we showed that the joint supervised learning algorithm provides a significant improvement over EM feature extraction, and benefits from an EM-based initialization. Furthermore, we demonstrated a significant improvement in classification accuracy when decoding postural changes involved in the tilt-table experiment, using the multivariate switching dynamics of HR and BP time series. The technique developed in this work is also significant from a theoretical point of view, since we demonstrated that one may apply backpropagation-based learning algorithm to inference in dynamic Bayesian networks. This results in *dynamical features* that are predictive of outcome. Since the EM learning objective is the log likelihood of the unlabeled time series, it may learn artifacts and other features that are not relevant to classification. As expected, increasing the number of EM steps in the simulation study (where high amplitude artifacts were randomly inserted into all time series) did not improve the discriminative performance, even though we observed a significant increase in training log likelihood. Notably, the proposed EM-based initialization step is qualitatively similar to the unsupervised learning step used for training Deep Belief Networks (DBN), where "unsupervised pre-training" is known to significantly improve the predictive performance of discriminative neural networks [10].

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

As with any learning algorithm, one must address potential problems with overfitting to the training data. Although, in the case of the simulated data we did not observe any overfitting (see Fig. 3(a)), the tilt-table results shown in Fig. 3(b) exhibits a slight performance decrease in two out of ten testing folds. We believe this is due to the small number of time series within the tilt-table cohort (10 time series versus 200 time series in the simulation study). Future work should involve exploring the full potential (and possible shortcomings) of the proposed learning algorithm for extracting dynamical features in large physiological cohort time series, with the goal of characterizing and predicting patient state and outcomes.

REFERENCES

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

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

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

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

[5] S. Nemati, L.-w. H. Lehman, R. P. Adams, and A. Malhotra, "Discovering shared cardiovascular dynamics within a patient cohort," in *Engineering in Medicine and Biology Society (EMBC), 2012 Annual International Conference of the IEEE*. IEEE, 2012, pp. 6526–6529.

[6] L.-w. H. Lehman, S. Nemati, R. P. Adams, and R. G. Mark, "Discovering shared dynamics in physiological signals: Application to patient monitoring in icu," in *Engineering in Medicine and Biology Society (EMBC), 2012 Annual International Conference of the IEEE*. IEEE, 2012, pp. 5939–5942.

[7] T. Heldt, M. B. Oefinger, M. Hoshiyama, and R. G. Mark, "Circulatory response to passive and active changes in posture," *Computers in Cardiology*, vol. 30, pp. 263–266, 2003, circulatory response to passive and active changes in posture.

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

[9] S. Nemati, "Identifying evolving multivariate dynamics in individual and cohort time-series, with application to physiological control systems," Ph.D. dissertation, MIT, December 2012.

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