Tracking Progression of Patient State of Health in Critical Care Using Inferred Shared Dynamics in Physiological Time Series

Li-wei H. Lehman[1], Shamim Nemati[2], Ryan P. Adams[2], George Moody[1], Atul Malhotra[3], and Roger G. Mark[1]

Abstract— Physiologic systems generate complex dynamics in their output signals that reflect the changing state of the underlying control systems. In this work, we used a switching vector autoregressive (switching VAR) framework to systematically learn and identify a collection of vital sign dynamics, which can possibly be recurrent within the same patient and shared across the entire cohort. We show that these dynamical behaviors can be used to characterize and elucidate the progression of patients' states of health over time. Using the mean arterial blood pressure time series of 337 ICU patients during the first 24 hours of their ICU stays, we demonstrated that the learned dynamics from as early as the first 8 hours of patients' ICU stays can achieve similar hospital mortality prediction performance as the well-known SAPS-I acuity scores, suggesting that the discovered latent dynamics structure may yield more timely insights into the progression of a patient's state of health than the traditional snapshot-based acuity scores.

I. INTRODUCTION

Physiologic systems generate complex dynamics in their output signals that reflect the changing state of the underlying control systems [1], [2], [3]. For instance, time series of blood pressure (BP) can exhibit oscillations on the order of seconds, to minutes, to hours [4]. Dynamical structure of vital sign time series may contain characteristic signatures of certain physiological and pathological states [5].

In [6], [7], [8], we proposed a switching linear dynamical systems (SLDS) framework to model nonlinear cardiovascular dynamics. The premise of our approach is that although the underlying dynamical system may be nonlinear and nonstationary, the dynamics can be well approximated by a mixture of linear dynamical systems. We assume that there exists a collection of possible dynamic behaviors that are exhibited across the vital sign time series of a patient cohort, and that each patient may take on a subset of these behaviors and transition between them upon a change in the underlying cardiovascular control system. Given a patient cohort, we seek to discover such a collection of dynamic behaviors or modes, and model each time series as switching between these dynamic modes. To formalize these objectives, we employed the switching vector autoregressive framework [9].

Given a collection of related time series from a cohort, the technique allows for simultaneous learning of the underlying shared dynamics, and identification of time series segments that follow a "similar" dynamic.

The goal of this current work is to investigate the utility of the proposed SLDS framework in tracking the progression of patients' states of health over time. We hypothesize that changes in the dynamical patterns of the patient's vital signs may be reflective of deterioration or recovery of the underlying cardiovascular control system, and that patients with different trajectories in the progression of their states of health also exhibit different evolution in the dynamical structure of their vital sign time series. Thus, by monitoring the changes in the dynamic structure of a patient's vital sign over time, one can obtain a more accurate and up-to-date assessment of a patient's state of health.

To test our hypothesis, we applied the proposed SLDS framework to model the blood pressure dynamics of an ICU patient cohort during the first 24 hours of their ICU stays. We compared the distribution of shared dynamic modes between non-survivors (patients who expired before discharge from the hospital) and survivors (all others), and examined the evolution of the learned dynamic structure of these two groups of patients over a 24-hour period. Next, we investigated the prediction performance of the dynamic modes in tracking patients' mortality risks over time. Finally, we showed example time series with the corresponding estimated mortality risks based on the inferred dynamical modes.

II. MATERIALS AND METHODS

This section describes the utilized dataset, as well as the technique of switching vector autoregressive process for discovery of shared dynamics among patients.

A. Dataset

Minute-by-minute mean arterial blood pressure (MAP) measurements of MIMIC II [10] adult ICU patients during the first 24 hours of their ICU stays were extracted. The analysis in this paper was restricted to 337 patients with day 1 SAPS-I scores [11] (from MIMIC II V25) and with at least 18 hours of blood pressure data during the first 24 hours after ICU admission. Hospital mortality of this cohort is 14%. The median SAPS-I score for this cohort is 16. Distributions of the 337 patients in care units are 20% coronary care unit (CCU), 45% Cardiac Surgery Recovery Unit (CSRU), 23%
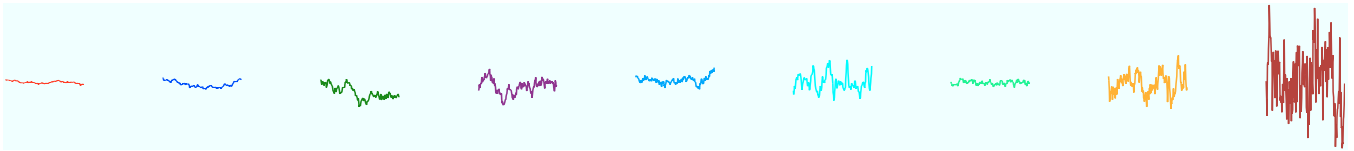
Fig. 1. Dynamic modes (left to right): 4 2 9 18 14 6 15 17 20. First two modes are "high-risk" modes. Next 5 modes are "low-risk" modes. The last two modes are neutral with "non-significant" associations with hospital mortality.

	Mode	Color	p-val	OR (95% CI)	AUC	AR Coef	Covariance
1	4	Red	0.0000	1.59 (1.28 1.98)	0.70	1.00, -0.00, -0.00	0.23
2	2	Blue	0.0005	1.71 (1.26 2.31)	0.64	0.92, 0.06, 0.01	0.63
3	9	Green	0.0049	0.28 (0.12 0.68)	0.66	1.03, -0.03, -0.01	4.46
4	18	Purple	0.0112	0.52 (0.31 0.86)	0.67	0.78, 0.06, 0.11	10.01
5	14	Light Blue	0.0091	0.58 (0.39 0.87)	0.64	0.67, 0.15, 0.16	3.69
6	6	Cyan	0.0286	0.31 (0.11 0.88)	0.59	1.48, -0.65, 0.07	9.32
7	15	Light Green	0.0260	0.23 (0.06 0.84)	0.62	0.79, -0.01, 0.00	2.22
8	17	Orange	0.7467	0.92 (0.55 1.54)	0.57	0.90, -0.11, 0.09	45.23
9	20	Brown	0.5939	1.74 (0.23 13.33)	0.54	0.56, -0.16, 0.27	627.77

TABLE I

DISCOVERED DYNAMIC MODES AND THEIR ASSOCIATIONS WITH HOSPITAL MORTALITY.

Medical Intensive Care Unit (MICU), 12% Surgical Intensive Care Unit (SICU).

B. Switching Vector Autoregressive Modeling of Cohort Time Series

Our approach to discovery of shared dynamics among patients was based on the switching vector autoregressive (VAR) model or the autoregressive HMMs (AR-HMM) technique [9]. We assume that there exists a library of possible dynamic behaviors or modes; a set of multivariate autoregressive model parameters $\{A_p, \ p = 1 \cdots P\}_{k=1}^K, \{Q\}_{k=1}^K$, where K is the number of dynamic modes, and for each mode k, $A_p^{(k)}, \ p = 1 \cdots P$ are the AR coefficient matrices (corresponding to each of the P lags) and $Q^{(k)}$ is the corresponding noise covariance matrix. Let y_t be the observation at time t, and z_t be an indicator variable indicating the active dynamic mode at time t. Following these definitions, an AR-HMM is defined by:

$$y_t = \sum_{p=1}^{P} A_p^{(z_t)} y_{t-p} + Q^{(z_t)}. \tag{1}$$

A collection of related time series can be modeled as switching between these dynamic behaviors which describe a locally coherent linear model that persist over a segment of time. We modeled minute-by-minute MAP time series as a switching AR(3) process with 20 dynamic modes.

C. Evaluation Methods and Statistical Analysis

Patients were divided into 10 training/test sets. Ten switching VAR models were learned, one for each training set. The mode assignment of time series for patients in the test set was inferred based on the model learned from the corresponding training set.

For each patient, we used the proportion of time a patient spent in each of the dynamic modes ("mode proportions"

from now on) to construct a feature vector for predicting a patient's underlying "state of health". Specifically, mode proportion is defined as the fraction of time patients stay within the different modes over a fixed time interval (6 or 24 hours in this work).

Univariate logistic regressions were performed to find associations between mode proportions and hospital mortality. For each mode, we report its p value, and odds ratios (OR, with 95% confidence interval) in mortality risks as per 10% increase in mode proportions; AUCs from the univariate analysis were reported as a measure of model fit.

Regularized logistic regression was used to predict hospital mortality based on mode proportions of patients in the test set; the regression weights were learned based on mode proportions (over the entire 24-hour period) of patients in the training set. To investigate the predictive value of dynamic modes over time during the first 24 hours after ICU admission, hourly risk scores were computed for each patient in the test set. Risk scores, computed as the probability of death from the logistical function, were based on mode proportions from a six-hour sliding window by stride of one hour over the entire 24 hour period. All hourly risk scores in this study were adjusted with care unit information. AUCs from 10 fold cross validations were reported as mean AUCs with standard deviations or median AUCs with the interquartile range.

III. RESULTS

A. Association between the Dynamic Mode Proportions and Hospital Mortality

We used univariate logistic regression analysis to test the association between the proportion of time patients spent in each dynamic mode (20 in total) during the first 24-hours in the ICU and the outcome (hospital mortality). Our results indicate that dynamic modes 4 and 2 are significant

($P < 0.0001$, $P < 0.001$) "high-risk" modes in which increased proportions of these modes are associated with higher hospital mortality with odds ratios 1.59 (1.28 1.98) and 1.71 (1.26 2.31) respectively.

Dynamic modes 9, 18, 14, 6, 15 are "low-risk" modes in which increasing proportions of these modes are significantly ($P < 0.05$) associated with a decreased risk of hospital mortality, with odds ratios less than one. See Figure 1 for illustrations of these dynamic modes. Table I lists the AR coefficients and covariance of the two high-risk and five low-risk dynamic modes, as well as their respective associations with hospital mortality. As a reference, we also show two modes (17 and 20) which did not have any significant associations with the hospital mortality. Note that the high-risk modes appear to correspond to less variability in their dynamics.

B. Dynamic Mode Proportions Over Time: Survivors vs. Non-Survivors

This section examines the changes in the dynamical structure of patients' blood pressure time series over the first 24 hours of patients' ICU stays. Figure 2 compares the survivors vs. the non-survivors in their proportions of time spent in the high-risk (modes 4 and 2) vs. the low-risk dynamic modes (modes 9, 18, 14, 6, 15) over the 24-hour period. Note that as time progressed, the non-survivors had an increasing trend in the proportion of time in the high-risk modes.

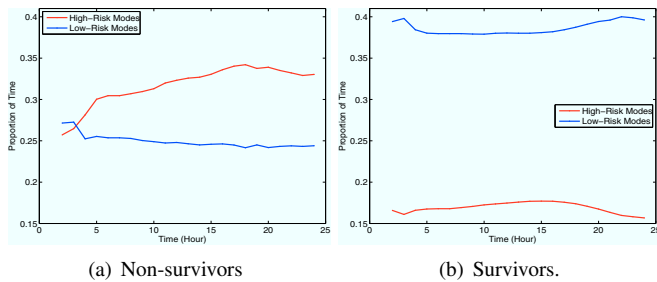(a) Non-survivors   (b) Survivors.

Fig. 2.   Comparison of (a) the non-survivors vs. (b) the survivors in the proportions of time each group spent in the high-risk dynamics vs. the low-risk dynamics during the first 24-hour period in the ICU. Plot shows the population mean mode proportions during each time interval. As time progressed, the non-survivor group had a significant increase in the proportion of time in the high-risk dynamic modes.

C. Prediction Performance Over Time

This section investigates the utility of the dynamic modes in tracking patients' mortality risks over time. Figure 3 displays the mortality prediction performance using mode proportions and the care unit information. The x-axis shows the time as the number of hours after ICU admission. AUC was computed at an hourly interval based on an exponential weighted moving average of the current and previous hourly risk scores. Note that the predictive performance of the dynamic modes increases as time progresses over the 24-hour period. The learned dynamics from as early as the first 8 hours of patients' ICU stays achieved similar hospital mortality prediction performance as the well-known SAPS-I
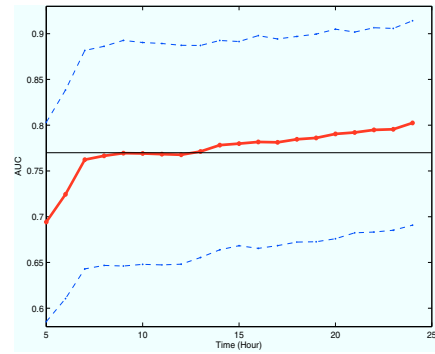

Fig. 3.   Mortality prediction performance over time using the hourly updated risk scores (adjusted with care unit information). Plot shows mean AUC (solid red line) with standard deviation (blue dashed line). Baseline performance (black line), using SAPS-I adjusted with care unit information, achieved a mean AUC of 0.77 ($\pm$ 0.07). SAPS-I alone achieved an average AUC of 0.65 ($\pm$ 0.05).

acuity scores. The mean and median AUCs of our technique using mode proportions and care unit information were 0.80 ($\pm$ 0.12) and 0.83 (interquartile range 0.70, 0.89) respectively at the end of the 24-hour period .

IV. EXAMPLE TIME SERIES OF PATIENTS WITH ESTIMATED MORTALITY RISKS OVER TIME

Figure 4 shows blood pressure time series for six patients. Panel (a) shows 3 patients with the highest risk scores (within the test set) at the end of the 24-hour period; panel (b) shows 3 patients with the largest decrease in their risk scores (within the same test set) during their first day in the ICU. All six patients were from the same test set, with mode assignment inferred based on dynamic modes learned from the corresponding training set. See Figure 1 and Table 1 for an illustration and description of the example learned dynamics.

Note that as time progresses, patients in panel (a) tend to spend more time in the "high-risk" dynamic modes (mode 2 in blue, mode 4 in red); their estimated mortality risks rise accordingly over time. In contrast, panel (b) patients show a decreasing trend in mortality risks as they transition to lower-risk dynamic modes over time.

V. DISCUSSION AND CONCLUSIONS

The goal of this current work is to investigate the utility of the proposed SLDS framework in tracking the progression of patients' states of health over time. We applied the framework to model the blood pressure dynamics of an ICU patient cohort during the first 24 hours of their ICU stays. We demonstrated that patients who did not survive the hospital stays exhibited different evolution in their vital sign dynamics than those who survived. These results support our hypothesis that, as patients improve or deteriorate in their states of health, the distribution of their vital sign dynamics also become increasingly different.

We evaluated the utility of our framework in continuous monitoring of patients' mortality risks on an hourly basis over the first 24-hour period. We showed that the discovered
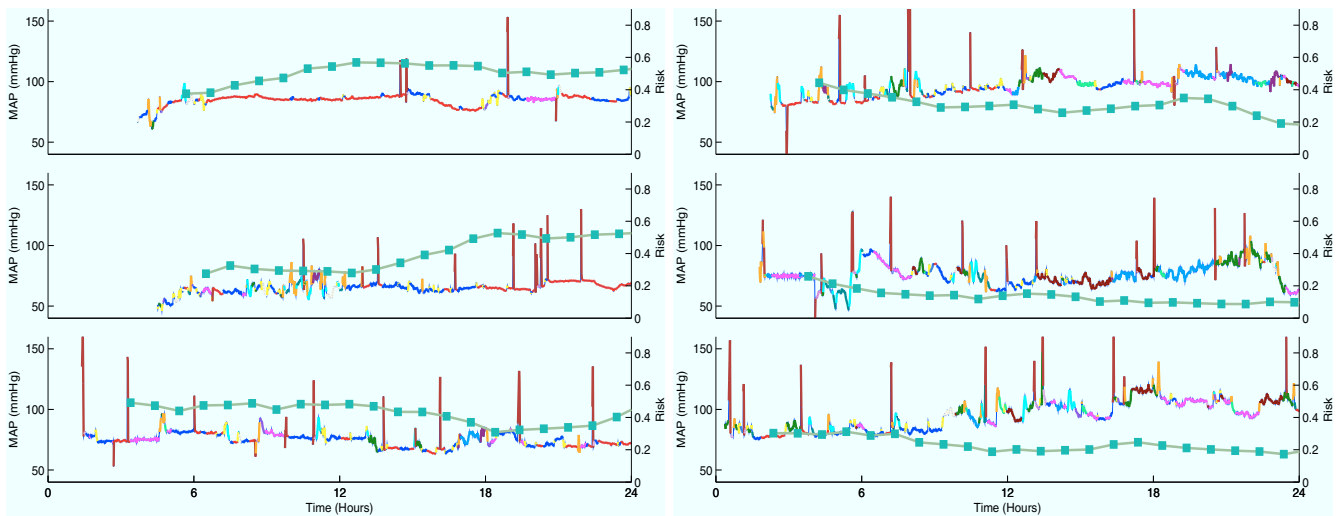
(a) Patients with the highest ending risk scores at the end of the first day ICU stay. All three patients were from MICU. All three patients deceased in the hospital. Hospital length-of-stay: 21, 2, 9 days.

(b) Patients with the largest decrease in their risk scores during their first day ICU stay. Care units: MICU, CCU, MICU. All three patients survived the hospital stay. Hospital length-of-stay: 46, 8, 9 days.

Fig. 4. Mean arterial blood pressure of six patients during the first 24-hours in the ICU. Samples color-coded by their mode assignment (Red - Mode 4; Blue - Mode 2; Green - Mode 9; Purple - Mode 18; Light Blue - Mode 14; Cyan - Mode 6; Light Green - Mode 15; Orange - Mode 17; Brown - Mode 20). Mortality risks were estimated every hour based on the inferred blood pressure dynamical modes in the previous six-hour period; estimated risks are plotted as dark green lines with scale indicated by y-axes on right-side of each graph. Blood pressure measurements plotted in original units (before de-trending).

dynamics can be used for prediction and tracking of a patient's propensity to survive hospital stays over time. Interestingly, the learned blood pressure dynamics from as early as the first 8 hours of patients' ICU stays can achieve similar hospital mortality prediction performance as the well-known SAPS-I acuity score [11], which included 14 lab measurements and demographics information. Further, the mortality prediction accuracy using the dynamics improved over time. Recent studies [4] suggest that therapeutic interventions not only should aim at maintaining the mean BP within an acceptable range, but also should direct the patient trajectory towards healthy dynamical regimes with enhanced variability. Our results suggest that continuous monitoring of the vital sign dynamics may provide valuable insights to the underlying cardiovascular control system and the disease progression of ICU patients, and could potentially be used to generate real-time predictive alerts to guide therapeutic interventions, and enable efficient and timely allocation of resources to high-risk patients.

Current and ongoing works involve combining the switching linear dynamical system framework with all available clinical data (lab tests, medication records, nursing notes, etc) in generating predictive alerts. Future work remains to investigate whether continuous monitoring based on dynamic modes can alert clinicians to deteriorating conditions in patients at an earlier stage than existing monitoring techniques, and whether such alerts lead to improved patient care and outcome.

ACKNOWLEDGMENT

REFERENCES

[1] 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, pp. 55-67, Jul 2011.

[2] M. Costa, A. L. Goldberger, and C. K. Peng, "Multiscale entropy analysis of complex physiologic time series," Phys Rev Lett, vol. 89, p. 068102, Aug 5 2002.

[3] P. C. Ivanov, L. A. Amaral, A. L. Goldberger, S. Havlin, M. G. Rosenblum, Z. R. Struzik, and H. E. Stanley, "Multifractality in human heartbeat dynamics," Nature, vol. 399, pp. 461-5, Jun 3 1999.

[4] G. Parati, J. E. Ochoa, C. Lombardi, and G. Bilo. "Assessment and management of blood-pressure variability," Nature Reviews Cardiology 10, 143-155, March 2013.

[5] A. L. Goldberger, L. A. Amaral, J. M. Hausdorff, P. Ivanov, C. K. Peng, and H. E. Stanley, "Fractal dynamics in physiology: alterations with disease and aging," Proc Natl Acad Sci U S A, vol. 99 Suppl 1, pp. 2466-72, Feb 19 2002.

[6] S. Nemati, L. Lehman, RP Adams, A Malhotra. "Discovering shared cardiovascular dynamics within a patient cohort". Proc 34th IEEE EMBS, 2012.

[7] L. Lehman, S. Nemati, RP Adams, R Mark. "Discovering shared dynamics in physiological signals: Application to patient monitoring in ICU". Proc 34th IEEE EMBS, 2012.

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

[9] K. P. Murphy, Switching Kalman Filter, Compaq Cambridge Research Laboratory, Tech. Rep. 98-10., 1998, cambridge, MA.

[10] M. Saeed, M. Villarroel, A. Reisner, G. Clifford, L. Lehman, G. Moody, T. Heldt, T.H. Kyaw, B. Moody, and R.G. Mark, "Multi-parameter intelligent monitoring in intensive care II (MIMIC-II): A public-access ICU database," Crit Care Med, 39(5), 2011.

[11] J.R. Le Gall, P. Loirat, A. Alperovitch, et al. "A simplified acute physiology score for ICU patients," Critical Care Medicine, 12(11), 975-977, 1984.