

Individualized Risk Prognosis for Critical Care Patients: A Multi-task Gaussian Process Model

Ahmed M. Alaa¹, Jinsung Yoon¹, Scott Hu², and Mihaela van der Schaar^{3,4,1}

¹ Electrical Engineering Department, University of California, Los Angeles, USA,
{ahmedmalaa, jsyoon0823}@ucla.edu,

² David Geffen School of Medicine, University of California, Los Angeles, USA,
scotthu@mednet.ucla.edu,

³ Department of Engineering Science, University of Oxford, UK,
mihaela.vanderschaar@eng.ox.ac.uk,

⁴ Alan Turing Institute, London, UK

Abstract. We report the development and validation of a *data-driven* real-time risk score that provides timely assessments for the clinical acuity of ward patients based on their temporal lab tests and vital signs, which allows for timely intensive care unit (ICU) admissions. Unlike the existing risk scoring technologies, the proposed score is “individualized”—it uses the *electronic health record* (EHR) data to cluster the patients based on their static covariates into *subcohorts* of “similar” patients, and then learns a separate temporal, non-stationary multi-task *Gaussian Process* (GP) model that captures the physiology of every subcohort. Experiments conducted on data from a heterogeneous cohort of 6,094 patients admitted to the Ronald Reagan UCLA medical center show that our risk score significantly outperforms the state-of-the-art risk scoring technologies, such as the *Rothman index* and *MEWS*, in terms of timeliness, true positive rate (TPR), and positive predictive value (PPV). In particular, the proposed score increases the AUC with 20% and 38% as compared to Rothman index and MEWS respectively, and can predict ICU admissions 8 hours before clinicians at a PPV of 35% and a TPR of 50%. Moreover, we show that the proposed risk score allows for better decisions on when to discharge clinically stable patients from the ward, thereby improving the efficiency of hospital resource utilization.

Keywords: Critical Care Prognostication, Gaussian Process, Intensive Care Unit, Personalized Medicine.

1 Introduction

Prognostic risk assessment models that quantify the acuity of critical care patients in real-time can inform vital and delay-critical clinical decision-making [1]. Unanticipated adverse events such as mortality, cardiopulmonary arrest, or intensive care unit (ICU) transfer are often preceded by disorders in a patient’s physiological parameters [2]. Timely prediction of such events can be carried out

by continuously quantifying the patient’s acuity using evidence in her physiological parameters, and hence assessing her risk for a specific event by prompting a real-time “risk score” that can be tracked by clinicians.

Current clinical practice in most hospitals and healthcare facilities rely on two categories of risk scoring technologies. The first category comprises *early-warning systems* (EWS), such as MEWS [3], which hinge on *expert-based* models for triggering transfer to the ICU. A major drawback of “expert-based” scores is that they are not subject to any rigorous, objective validation. Recent systematic reviews have shown that EWS-based scores only marginally improve patient outcomes while substantially increasing clinician and nursing workloads, leading to alarm fatigue and inefficient resource utilization [4–7]. The second category of risk scores are based on data-intensive regression models that are built using the electronic health record (EHR) data. The most notable technology in this category is the Rothman index [8], which is currently deployed in more than 70 hospitals in the US (including the Houston Methodist hospital in Texas, and Yale-New Haven hospital in Connecticut) [9], and was recently shown to be superior to MEWS-based models in terms of false alarm rates [10].

While the Rothman index offers a significant performance improvement over MEWS, it suffers from 2 major drawbacks. First, it adopts a “one-size-fits-all” risk scoring model that ignores the individual traits of the monitored patients. Second, it ignores the temporal aspect of the physiological data as it computes the patient’s risk score at a particular moment using the patient’s vital signs at that moment, ignoring her temporal physiological trajectory (see Figure 1 in [8]). To that end, we report the development and validation of a novel data-driven real-time risk score that addresses these drawbacks and provides a significant performance improvement over all existing clinical risk scoring technologies. The proposed risk score is a numeric value between 0 and 1 that corresponds to the patient’s risk for clinical deterioration, and is computed and updated in real-time by aggregating two types of a monitored patient’s information:

1. **Static Admission Information:** this includes all the static information gathered about the patient upon her hospitalization and remain fixed during her stay in the ward (e.g. age, gender, race, ICD-9 code, diagnosis, etc).
2. **Time-Varying Physiological Information:** this includes all the physiological parameters (vital signs and lab tests) that are repeatedly gathered for the patient during her stay in the ward (e.g. systolic and diastolic blood pressure, O_2 saturation, respiratory rate, Glucose, Glasgow coma scale score, Creatinine, etc).

Our risk score models a patient’s entire temporal physiological trajectory via a non-stationary multi-task *Gaussian Process* (GP) model, which captures irregularly sampled and temporally correlated physiological data [11]. The model parameters are learned in a data-driven fashion: we use the EHR data in order to fit the multi-task GP *hyper-parameters* for clinically stable patients (patients who were recorded in the EHR as discharged from the ward), and fit a different set of hyper-parameters for clinically deteriorating patients (patients who were

Table 1. Physiological data and admission information associated with each patient in the cohort under study.

Time-Varying Physiological Data		
Vital Signs	Lab Tests	Static Admission Information
Diastolic blood pressure	Glucose	Age
Eye opening	Urea Nitrogen	Admission floor
Glasgow coma score	White blood cell	Gender
Heart rate	Creatinine	Stem cell transplant
Respiratory rate	Hemoglobin	ICD-9 code
Temperature	Platelet Count	Transfer status
O ₂ Device Assistance	Potassium	
O ₂ Saturation	Sodium	
Best motor response	Total CO ₂	
Best verbal response	Chloride	
Systolic blood pressure		

recorded in the EHR as transferred to the ICU). The patient’s risk score is computed as the optimal test statistic of a *sequential hypothesis test* that tests the hypothesis that the patient is clinically deteriorating given a sequence of physiological measurements [12]. Following the newly emerging concepts of *precision medicine*, we ensure that our risk scoring procedure is tailored to the individual’s traits by introducing *latent phenotype* variables, where a phenotype represents a distinct way in which a patient manifests her clinical status [13]. Using unsupervised learning, we discover the number of patient phenotypes from the EHR data, learn the association between a patient’s static admission information with her phenotype, and calibrate a separate model for every phenotype.

2 Study Subjects

We conducted our experiments using an EHR dataset for a heterogeneous cohort of 6,094 patients admitted to the Ronald Reagan UCLA medical center in a period that spans 3 years (March 2013 to March 2016). The patients’ cohort is quite heterogeneous; we considered admissions to all units in the medical center, including the cardiac observation unit, cardiothoracic unit, hematology and stem cell transplant unit and the liver transplant service. The cohort comprised patients with a wide variety of ICD-9 codes and medical conditions, including leukemia, hypertension, sepsis, abdomen, pneumonia, and renal failure.

Every patient record in the dataset is associated with the time-varying and static information listed in Table 1. The time-varying physiological measurements are collected over irregularly spaced time intervals (usually ranging from 1 to 4 hours); for each physiological time series, we have access to the times at which each value was gathered. The patients’ length of stay in the wards ranged from 4 hours to 2700 hours. Every patient in the cohort is associated with a label that indicates whether the patient was discharged home or transferred to

the ICU, and hence we know which patients were clinically stable and which ones were clinically deteriorating. Around 9% of the patients in the cohort were clinically deteriorating and experienced an unanticipated ICU transfer. We excluded all patients for whom a routine ICU transfer was preordained since the objective of the proposed risk score is to predict unanticipated ICU admissions.

3 The Proposed Individualized Risk Score

3.1 Notations and Definitions

Let $X_i(t) = [X_{i1}(t), X_{i2}(t), \dots, X_{iD}(t)]^T$ be a D -dimensional stochastic process representing the D time-varying physiological streams for patient i . In our dataset, $D = 21$, i.e. the number of lab tests and vital signs listed in Table 1. The vital signs and lab tests of patient i are gathered at arbitrary time instances $\{t_{dj}^i\}_{d=1, j=1}^{D, M_d}$ (where $t = 0$ is the time at which the patient is admitted to the ward), where M_d is the total number of measurements of the d^{th} vital sign (or lab test) that were gathered during the patient's stay in the ward. Thus, the set of all observations of the physiological data that clinicians collect for a specific patient during her stay in the ward is given by $\{X_i(t_{dj}^i)\}_{d=1, j=1}^{D, M_d}$, and we will refer to the realizations of these variables as $\{x_{dj}^i, t_{dj}^i\}_{i, d, j}$.

We define the S -dimensional random vector Y_i as patient i 's static admission information, i.e. in our dataset, $S = 6$ as indicated in Table 1. We denote the realizations of patient i 's static information as $Y_i = y_i$. Thus, the set of all information associated with a patient can be gathered in a set $\{y_i, x_{dj}^i, t_{dj}^i\}_{i, d, j}$.

3.2 Risk Scoring as a Sequential Hypothesis Test

Let $V_i \in \{0, 1\}$ be a binary *latent* variable that corresponds to patient i 's true clinical status, where $V_i = 0$ stands for a stable clinical status, and $V_i = 1$ stands for a deteriorating clinical status. We assume that V_i is a fixed *latent class* that determines the physiological model, i.e. V_i is drawn randomly for patient i at admission time and stays fixed over the patient's stay in the ward. In the EHR dataset, the value of V_i is revealed at the end of every physiological stream, where $V_i = 1$ if the patient is admitted to the ICU, and $V_i = 0$ if the patient is discharged home.

Since physiological streams manifest the patients' clinical statuses, it is natural to assume that the conditional distributions of $\{x_{dj}^i, t_{dj}^i\}_{i, d, j}$ given $V_i = 0$ differs from that of $\{x_{dj}^i, t_{dj}^i\}_{i, d, j}$ given $V_i = 1$. Our conception of the risk score can be described as follows. During the patient's stay in the ward, we are confronted with two hypotheses: the null hypothesis \mathcal{H}_0 corresponds to the hypothesis that the patient is clinically stable, whereas the alternative hypothesis \mathcal{H}_1 corresponds to the hypothesis that the patient is clinically deteriorating, i.e.

$$V_i = \begin{cases} 0 : \mathcal{H}_0 \text{ (clinically stable patient)}, \\ 1 : \mathcal{H}_1 \text{ (clinically deteriorating patient)}. \end{cases} \quad (1)$$

Thus, the prognosis problem is equivalent to a *sequential hypothesis test* [12], i.e. the clinicians need to reject one of the hypotheses at some point of time after observing a series of physiological measurements. Hence, we view the patient’s risk score as the test statistic of a sequential hypothesis test: patient i ’s risk score at time t , which we denote as $R_i(t) \in [0, 1]$, is the posterior probability of hypothesis \mathcal{H}_1 given the observations $\{y_i, x_{dj}^i, t_{dj}^i \leq t\}_{i,d,j}$, i.e. $R_i(t) = \mathbb{P}(\mathcal{H}_1 \mid \{y_i, x_{dj}^i, t_{dj}^i \leq t\}_{i,d,j})$. Using Bayes’ rule we have that

$$R_i(t) = \frac{\mathbb{P}(\{y_i, x_{dj}^i, t_{dj}^i \leq t\}_{i,d,j} \mid \mathcal{H}_1) \cdot \mathbb{P}(\mathcal{H}_1)}{\sum_{v \in \{0,1\}} \mathbb{P}(\{y_i, x_{dj}^i, t_{dj}^i \leq t\}_{i,d,j} \mid \mathcal{H}_v) \cdot \mathbb{P}(\mathcal{H}_v)}, \quad (2)$$

where $\mathbb{P}(\mathcal{H}_1) = \mathbb{P}(V_i = 1)$ is the prior probability of a patient in the ward being admitted to the ICU (i.e. the rate of ICU admissions, which is 9% in our dataset). In order to be able to compute the risk score in (2), we specify the conditional distributions $\mathbb{P}(\{y_i, x_{dj}^i, t_{dj}^i \leq t\}_{i,d,j} \mid \mathcal{H}_v)$, $v \in \{0, 1\}$ in the next Subsection.

3.3 The Non-stationary Multi-task GP Model

We adopt a multi-task GP model as the distributional specification for the continuous-time process $X_i(t)$; our choice for a multi-task GP is motivated by its ability to capture irregularly sampled, multi-variate time series data [11]. We specify a different set of GP hyper-parameters for the physiological data generated under $V_i = 0$ and those generated under $V_i = 1$, i.e.

$$X_i(t) \mid V_i = v \sim \mathcal{GP}(\Theta_v), \quad (3)$$

where Θ_v is the hyper-parameter set of the GP model for patients with $V_i = v$.

Since hospitalized patients are subject to external clinical interventions, and are likely to exhibit a progression of subsequent “phases” of clinical stability or deterioration, a stationary covariance kernel, such as the one defined in [11], would do not suffice to describe the patients’ entire physiological trajectory. This motivates a non-stationary model for $X(t)$ in which the time domain is divided into a sequence of K *epochs*, where the k^{th} epoch for patient i has a duration of T_k^i and is described by a *locally stationary* covariance kernel. This is achieved by assigning a different set of mean and stationary covariance hyper-parameters $\Theta_{v,k}$ for every epoch $1 \leq k \leq K$, and assuming that the physiological data in different epochs are independent. The duration of epoch k for patient i , T_k^i , is an integer number of hours that is drawn from a *negative binomial distribution* $f_{kv}(T \mid \lambda_{kv})$ with a parameter λ_{kv} . Since patients arrive at the hospital ward at random time instances, at which the clinical status is unknown, we define $\bar{k}_i \in \{1, 2, \dots, K\}$ as a latent initial epoch index for patient i , which we assume to be drawn from a *multinomial* distribution. Thus, the clinicians observe physiological measurements drawn from a process with the underlying epoch index sequence $\{\bar{k}_i, \bar{k}_i + 1, \dots, K\}$, with random epoch durations $\{T_{\bar{k}_i}^i, T_{\bar{k}_i+1}^i, \dots, T_K^i\}$, and for which the GP epoch-specific hyper-parameters

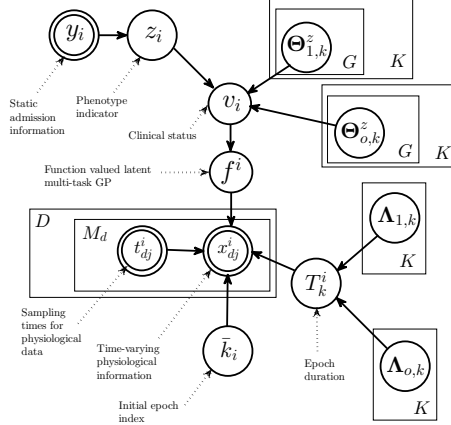


Fig. 1. Graphical depiction for the proposed physiological model.

are $\{\Theta_{v,\bar{k}_i}, \Theta_{v,\bar{k}_i+1}, \dots, \Theta_{v,K}\}$. Note that we assume that all the patients' epoch sequences end with epoch K , this ensures that all the physiological time series of all patients in the cohort are temporally *aligned*, which is essential for proper learning of the model parameters.

The GP hyper-parameters $\Theta_{v,k}$ for clinical state v and epoch k comprise a constant mean functions $m_{v,k}$ and a *squared exponential* covariance kernel with an *intrinsic correlation model* for the correlations between the different vital signs and lab tests [14]. That is, the covariance kernel $K_{v,k}(u, w, t, t')$ which quantifies the correlation between the physiological measurements $X_{iu}(t)$ and $X_{iw}(t')$ is given by

$$K_{v,k}(u, w, t, t') = \Sigma_{v,k}(u, w) k_{v,k}(t, t'), \quad (4)$$

if t and t' belong to the same epoch, and $k(u, v, t, t') = 0$ otherwise. In (4), $\Sigma_{v,k}$ is a positive semi-definite correlation matrix, and $k_{v,k}(t, t')$ is a squared exponential covariance kernel given as follows

$$k_{v,k}(t, t') = \exp\left(-\frac{1}{2\ell_{v,k}^2} \|t - t'\|^2\right), \quad (5)$$

where $\ell_{v,k}$ is the *characteristic length scale* parameter of the GP. We denote the set of all GP hyper-parameters under clinical status v as $\Theta_v = \{\Theta_{v,1}, \dots, \Theta_{v,K}\}$, and the epoch duration parameters as $\Lambda_v = \{\lambda_{v,1}, \dots, \lambda_{v,K}\}$.

3.4 The Latent Phenotype Indicators

The model presented in Subsection 3.3 is a “one-size-fits-all” model that treats all patients in the same way since it does not incorporate the patients' baseline

static features; however, in reality, different patients manifest their clinical status in different ways depending on their traits— the different ways in which patients manifests their physiology are known as the *phenotypes* [13]. In this Subsection, we refine the model in Subsection 3.3 to ensure *individualization*, i.e. ensure that our physiological model is tailored to the individual traits, by introducing a *latent phenotype indicator variable* $Z_i \in \{1, 2, \dots, G\}$ as the phenotype to which patient i belongs. We assume that the phenotype indicator variable Z_i possesses the following properties: $Z_i \perp\!\!\!\perp V_i | Y_i$, and $V_i \perp\!\!\!\perp Y_i | Z_i$.

We assume that a separate GP model is associated with every phenotype, i.e. for phenotype $z \in \{1, 2, \dots, G\}$, the corresponding GP hyper-parameter set is Θ_v^z , and the epoch duration parameter set is Λ_v^z . The phenotype indicators are latent and hence we do not know to which phenotype patient i belongs upon admission. However, since a patient’s phenotype naturally depends on her age, race, diagnosis, etc, we can infer the patient’s phenotype by estimating the posterior probability of patient i ’s membership in phenotype z , which we denote as $\gamma_z(Y_i) = \mathbb{P}(Z_i = z | Y_i)$. Thus, the resulting physiological model is a *mixture model* that combines G instantiations of the model in Subsection 3.3 with weights that are proportional to the phenotype memberships $\{\gamma_1(Y_i), \dots, \gamma_G(Y_i)\}$. The model parameters Θ_v^z and Λ_v^z are estimated from the dataset using the standard expectation-maximization (EM) algorithm, and we use the *Bayesian information criterion* to select the number of phenotypes G and the number of epochs K .

Figure 1 depicts a graphical model for the patients’ physiological data. In our previous works in [15] and [16], we developed *ForecastICU*, a GP-based risk score that we have shown to be superior to both the Rothman index and MEWS. *ForecastICU* is a subset of the proposed model that does not consider individualization and non-stationarity; we compare the performance of the proposed risk model with *ForecastICU* in the next Section.

4 Results

We evaluated the prognostic utility of the proposed risk scoring algorithm by dividing the patient’s cohort into a training set (admissions between March 2013 and November 2015) and a testing set (admissions between November 2015 and March 2016), estimating the model parameters from the training set, and then emulating the ICU admission decisions on the testing set by setting a threshold on the risk score $R_i(t)$, above which patient i is identified as “clinically deteriorating”. The accuracy of such decisions are assessed via the following performance metrics: true positive rate (TPR), positive predictive value (PPV) and timeliness (i.e. the difference between the time of actual ICU admission as decided by clinicians and the time at which $R_i(t)$ exceeds the threshold). Using the Bayesian information criterion, we selected an instantiation of our model with 12 epochs and 4 phenotypes.

In Figure 2, we compare the ROC curve of the proposed risk score with those of the Rothman index and MEWS scores. In addition, we compare our risk score with the APACHE II and SOFA scores; both scores were originally developed

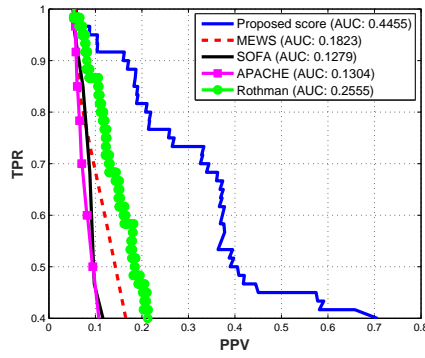


Fig. 2. ROC curve (TPR vs PPV).




Fig. 3. Timeliness curve.

Table 2. Performance comparisons for various machine learning algorithms ($p < 0.01$).

Algorithm	Proposed risk score	ForecastICU	Random Forest	Logistic regression	LASSO	RNN	HMM	MTGP	Rothman	MEWS	SOFA	APACHE
AUC (ICU admission)	0.45	0.39	0.36	0.27	0.26	0.29	0.32	0.3	0.25	0.18	0.13	0.14
AUC (Discharge at 0.01)	0.36	0.3	0.2	0.17	0.16	0.18	0.23	0.23	0.25	0.18	0.1	0.13
AUC (Discharge at 0.05)	0.34	0.25	0.17	0.15	0.13	0.17	0.19	0.2	0.24	0.18	0.1	0.13
AUC (Discharge at 0.2)	0.23	0.19	0.13	0.1	0.1	0.13	0.16	0.14	0.14	0.17	0.1	0.1

to predict mortality in the ICU but were recently shown to possess significant predictive power for predicting clinical deterioration in wards [17]. As we can see in Figure 2, the proposed risk score significantly outperforms all the other risk scores in terms of the AUC. In particular, the proposed risk score’s AUC is 20% greater than that of the Rothman index, the best performing clinical risk score. Since over 200,000 in-hospital cardiac arrests occur in the U.S. annually [18], those performance improvements gains correspond to thousands of lives saved each year. Moreover, as we can see in Figure 2, the proposed risk score offers a greater value for the PPV at any TPR value. This result has a great implication in clinical practice; the proposed risk score can ensure more confidence in its issued ICU alarms, which would mitigate alarm fatigue and enhance a hospital’s resource utilization [7]. The key behind the performance gains achieved by our risk score is that it considers the patient’s entire physiological trajectory and not just the current physiological measurements, thus it is not impulsively triggered

by instantaneous physiological data that may not be truly reflective of clinical deterioration, thereby reducing the rate of false alarms.

In Figure 3, we fix the TPR value at 50% and evaluate the timeliness of various risk score at different values of the PPV. The resulting curve, which we call the *timeliness curve*, illustrates the trade-off between the timeliness of the ICU alarms and the false alarm rates, i.e. the more quick a risk score issues ICU alarms, the more likely it will exhibit a false alarm. As we can see in Figure 3, the proposed risk score is always many hours ahead of all other scores for any value of the PPV, and can help predict ICU admissions many hours before a clinician would do at a reasonable PPV and TPR. For instance, the proposed risk score can predict ICU admissions 8 hours before clinicians for a PPV of 35% and TPR of 50%. By issuing prompt alarms for clinically deteriorating patients, the proposed risk score with can provide clinicians with a safety net to focus their attentions on patients who are more likely to deteriorate many hours before they exhibit severe decompensation, allowing for timely ICU admission and more efficient therapeutic interventions.

We also compared our risk score with other machine learning algorithms, including random forests, ForecastICU, logistic regression, recurrent neural networks (RNNs), hidden Markov models with Gaussian emissions (HMMs), multi-task GP regression (MTGPs) [11], and LASSO. Each of these algorithms is trained using a window of physiological measurements that precedes the ICU admission or discharge decision, and the size of such window is optimized separately for every algorithm. The AUC results for all the algorithms under consideration are listed in Table 2, and as we can see, the proposed outperforms all the competing algorithms, including our previous work in [15], which did not consider individualization. We also evaluated the AUC of our risk score and all competing algorithms when jointly issuing both ICU and discharge alarms. That is, we set a lower risk threshold, below which the patient is discharged home, and an upper threshold above which the patient is transferred to the ICU. We computed the AUC of all risk scores when fixing the lower risk threshold at values 0.01, 0.05 and 0.2, and sweeping the upper risk threshold from 0 to 1. In all cases, the proposed risk score outperforms all the other benchmarks. Thus, the proposed risk score can help making better utilization for hospital resources by discharging patients who are clinically stable from the ward in a timely manner.

5 Conclusions

In this study, a novel real-time risk score for actionable ICU prognostication is developed and validated. Unlike state-of-the-art risk models, the proposed risk score incorporates both the patients' non-stationary temporal physiological information and their individual baseline co-variables in order to accurately describe the patients' physiological trajectories. Experiments conducted on a cohort of 6,094 patients admitted to the Ronald Reagan UCLA medical center show that the proposed risk score is significantly more accurate than currently deployed risk scores and other machine learning algorithms.

References

1. Churpek MM, Yuen TC, Park SY, Gibbons R, Edelson DP. Using electronic health record data to develop and validate a prediction model for adverse outcomes on the wards. *Critical care medicine*. 2014;42(4):841.
2. Kaese J, Smith G, Prytherch D, Parr M, Flabouris A, Hillman K, et al. A comparison of antecedents to cardiac arrests, deaths and emergency intensive care admissions in Australia and New Zealand, and the United Kingdom the ACADEMIA study. *Resuscitation*. 2004;62(3):275–282.
3. Morgan R, Williams F, Wright M. An early warning scoring system for detecting developing critical illness. *Clin Intensive Care*. 1997;8(2):100.
4. Tsien CL, Fackler JC. Poor prognosis for existing monitors in the intensive care unit. *Critical care medicine*. 1997;25(4):614–619.
5. Cvach M. Monitor alarm fatigue: an integrative review. *Biomedical Instrumentation & Technology*. 2012;46(4):268–277.
6. Bliss JP, Dunn MC. Behavioural implications of alarm mistrust as a function of task workload. *Ergonomics*. 2000;43(9):1283–1300.
7. Subbe C, Kruger M, Rutherford P, Gemmel L. Validation of a modified Early Warning Score in medical admissions. *Qjm*. 2001;94(10):521–526.
8. Rothman MJ, Rothman SI, Beals J. Development and validation of a continuous measure of patient condition using the electronic medical record. *Journal of biomedical informatics*. 2013;46(5):837–848.
9. Landro L. Hospitals Find New Ways to Monitor Patients 24/7. *The Wall Street Journal*. 2015;.
10. Finlay GD, Rothman MJ, Smith RA. Measuring the modified early warning score and the Rothman index: advantages of utilizing the electronic medical record in an early warning system. *Journal of hospital medicine*. 2014;9(2):116–119.
11. Ghassemi M, Pimentel MA, Naumann T, Brennan T, Clifton DA, Szolovits P, et al. A Multivariate Timeseries Modeling Approach to Severity of Illness Assessment and Forecasting in ICU with Sparse, Heterogeneous Clinical Data. In: *AAAI*; 2015. p. 446–453.
12. Wald A. *Sequential analysis*. Courier Corporation; 1973.
13. Saria S, Goldenberg A. Subtyping: What it is and its role in precision medicine. *Intelligent Systems, IEEE*. 2015;30(4):70–75.
14. Bonilla EV, Chai KM, Williams C. Multi-task Gaussian process prediction. In: *Advances in neural information processing systems*; 2007. p. 153–160.
15. Yoon J, Alaa AM, Hu S, van der Schaar M. ForecastICU: A Prognostic Decision Support System for Timely Prediction of Intensive Care Unit Admission. In: *ICML*; 2016. p. 1680–1689.
16. Alaa AM, Yoon J, Hu S, van der Schaar M. A Semi-Markov Switching Linear Gaussian Model for Censored Physiological Data. *Workshop on Machine Learning for Health, NIPS*. 2016;.
17. Yu S, Leung S, Heo M, Soto GJ, Shah RT, Gunda S, et al. Comparison of risk prediction scoring systems for ward patients: a retrospective nested case-control study. *Critical Care*. 2014;18(3):1.
18. Merchant RM, Yang L, Becker LB, Berg RA, Nadkarni V, Nichol G, et al. Incidence of treated cardiac arrest in hospitalized patients in the United States. *Critical care medicine*. 2011;39(11):2401.