

---

# Learning from Clinical Judgments: Semi-Markov-Modulated Marked Hawkes Processes for Risk Prognosis

---

Ahmed M. Alaa,<sup>1</sup> Scott Hu,<sup>1</sup> Mihaela van der Schaar<sup>1 2 3</sup>

## Abstract

Critically ill patients in regular wards are vulnerable to unanticipated adverse events which require prompt transfer to the intensive care unit (ICU). To allow for accurate prognosis of deteriorating patients, we develop a novel continuous-time probabilistic model for a monitored patient’s temporal sequence of physiological data. Our model captures “*informatively sampled*” patient *episodes*: the clinicians’ decisions on when to observe a hospitalized patient’s vital signs and lab tests over time are represented by a *marked Hawkes process*, with intensity parameters that are modulated by the patient’s latent *clinical states*, and with observable physiological data (mark process) modeled as a *switching multi-task Gaussian process*. In addition, our model captures “*informatively censored*” patient episodes by representing the patient’s latent clinical states as an *absorbing semi-Markov jump process*. The model parameters are learned from offline patient episodes in the electronic health records via an EM-based algorithm. Experiments conducted on a cohort of patients admitted to a major medical center over a 3-year period show that risk prognosis based on our model significantly outperforms the currently deployed medical risk scores and other baseline machine learning algorithms.

## 1. Introduction

Hospitalized patients are vulnerable to a wide range of adverse events, including cardiopulmonary arrests (Kause et al., 2004; Hogan et al., 2012; Yoon et al., 2016), acute respiratory failures (Mokart, 2013), septic shocks (Henry et al., 2015), and post-operative complications (Clifton

et al., 2012). For a patient in a regular ward, the occurrence of any such event entails an unplanned transfer to an intensive care unit (ICU), the timing of which is a major determinant of the eventual outcome. Indeed, recent medical studies have confirmed that delayed transfer to the ICU is strongly correlated with morbidity and mortality (Mardini L, 2012; Mokart, 2013). The problem of delayed ICU transfer is enormous and acute: over 750,000 septic shocks and 200,000 cardiac arrests occur in the U.S. each year with mortality rates of 28.6% and 75% respectively (Merchant et al., 2011; Kumar et al., 2011). Fortunately, experts believe that much of these events could be prevented with accurate *prognosis* and early warning (Nguyen et al., 2007).

Motivated by the proliferation of electronic health records (EHRs) (currently available in more than 75% of hospitals in the U.S. (Charles et al., 2016)) we develop a *data-driven* real-time risk score that can promptly assess a hospitalized patient’s risk of clinical deterioration. Our risk score hinges on a novel continuous-time *semi-Markov-modulated marked Hawkes process* model for a monitored patient’s *episode*, i.e. the patient’s evolving (latent) clinical states and her corresponding (observed) physiological data. With the guidance of critical care experts, we conducted experiments on a dataset for a cohort of critically-ill patients admitted to a major academic medical center. Results show that our risk score offers significant gains in the accuracy (and timeliness) of predicting clinical deterioration; our risk score attains a 23% improvement in the Area Under Receiver Operating Characteristic (AUROC) as compared to the technology currently deployed in our medical center. Since it confers a significant prognostic value in subacute care in wards, the proposed risk score is currently being installed in our medical center as a replacement for the current technology.

The proposed probabilistic model (based on which our risk score is computed) captures a hospitalized patient’s entire *episode* as recorded in the EHR. A typical (critical care) patient episode comprises the time of her admission to the ward, the time of her admission to the ICU or discharge from the ward, and a temporal sequence of irregularly sampled physiological data that are collected during her stay in the ward (Johnson, 2016; Ghassemi et al., 2015). We model

---

<sup>\*</sup>Equal contribution <sup>1</sup>University of California, Los Angeles, US. <sup>2</sup>University of Oxford, UK. <sup>3</sup>Alan Turing Institute, UK.. Correspondence to: Ahmed M. Alaa <ahmedmalaa@ucla.edu>.

a patient’s episode as being driven by a latent *clinical state process*, which we represent as a *semi-Markov jump process* (Yu, 2010), describing the evolution of the patient’s “severity of illness” over time. All the observable physiological variables are modulated by this process: the times at which clinicians decide to observe the patient’s physiological data are drawn from a *Hawkes point process* (Hawkes & Oakes, 1974), the intensity of which is modulated by the patient’s clinical state process, whereas the observed physiological data is drawn from a *switching multi-task Gaussian process* with hyper-parameters that depend on the patient’s clinical state. The patient episode is thus a *marked Hawkes process*—with the physiological data serving as the *marks*—that is modulated by the patient’s clinical states. We provide a detailed description of the model in Section 3, and then propose an EM-based algorithm for learning its parameters from the EHR data in Section 4.

A distinctive feature of our model is its ability to incorporate informative *clinical judgments* into the generative process for the patient’s episode. The manifestation of informative clinical judgments in the EHR episodes is double-faceted: the patients’ episodes are both “*informatively sampled*” and “*informatively censored*”. Informative sampling results from the fact that clinicians decide to observe the patient’s physiological data more intensely if they believe that the patient is in a “bad” clinical state (Moskovitch et al., 2015; Qin & Shelton, 2015)—a belief that is based on either the clinician’s own assessment of the patient’s state, or the communication between the patient and the ward staff (Kyriacos et al., 2014). Informative censoring results from the clinicians’ decision on when to send the patient to the ICU or discharge her from the ward, which is indeed informative of the “clinical deterioration” or “clinical stability” onsets. In our model, informative censoring is taken into account by adopting an *absorbing* semi-Markov chain as a model for the patient’s latent clinical states; a patient’s risk score at any point of time is thus defined as the *probability of eventual absorption* in a “clinical deterioration” state.

**Related work:** Marked point processes have been recently used in a very different context to model *check-in* data (Du et al., 2016; Pan et al., 2016), but we are not aware of any attempts for their deployment in the medical context. Most of the previous works on risk prognosis for critical care patients viewed informative censoring as a “*surrogate label*” for a patient’s clinical deterioration, and hence used those labels to train a supervised (regression) model using the physiological data in a fixed-size time window before censoring. The supervised models used in the literature included logistic regression (Ho et al., 2012; Saria et al., 2010), SVMs (Wiens et al., 2012), Gaussian processes (Ghassemi et al., 2015; Yoon et al., 2016) and recurrent neural networks (Che et al., 2016). The main limitation of this approach is that, in addition to the fact that it gener-

ally does not deal with informatively sampled episodes, it does not model the entire patient’s physiological trajectory, and hence it does not accurately capture intermediate (subtle) deterioration stages that are indicative of future severe deterioration stages, which leads to a sluggish risk assessment and delayed ICU alarms.

Another strand of literature has focused on building probabilistic models, usually variants of Hidden Markov Models (HMMs), for the entire patient’s physiological trajectories; applications have ranged from disease progression modeling to neonatal sepsis prediction (Wang et al., 2014; Stanculescu et al., 2014). These models are not capable of dealing with irregularly-sampled data, do not deal with informatively sampled episodes, and are restricted to the Markovianity assumption which entails unrealistically memory-less clinical state transitions. In (Henry et al., 2015), a ranking algorithm was used to construct a risk score for sepsis shocks; however, the approach therein requires the clinicians to provide assessments that order the disease severity at different time instances— we typically do not have such data in the EHR for ward patients.

Various works in the medical literature have proposed “expert-based” medical risk scores for prognosis in hospital wards (Morgan et al., 1997; Parshuram et al., 2009), some of which are currently used in practice. The medical literature has also suggested the use of mortality risk scores that are normally used in the ICU, such as APACHE-II and SOFA, as risk scores for ward patients (Yu et al., 2014). However, recent systematic reviews have demonstrated the modest net clinical utility of all these scores (Cvach, 2012). More recently, a data-driven medical risk score based on a simple regression model, known as the *Rothman index*, has been developed and commercialized (Finlay et al., 2014; Rothman et al., 2013). The Rothman index is currently deployed in various major hospitals in the U.S. including our medical center; in Section 5, we show that our risk score significantly outperforms the Rothman index in terms of AUROC and timeliness of ICU admission alarms.

## 2. Structure of the EHR Data

The subacute care data in an EHR typically comprises a set of *episodes*; each episode is a sequence of vital signs and lab tests (physiological data) that have been gathered (by clinicians) for a hospitalized patient at irregularly spaced time instances during her stay in a ward. The episode starts at the time of admission to the ward, and is concluded by either an *unplanned admission to the ICU*, which means that the patient was *clinically deteriorating*, or a *discharge from the ward*, which means that the patient was *clinically stable*. We denote an EHR dataset  $\mathcal{D}$  that comprises the episodes for  $D$  patients as  $\mathcal{D} = \{\mathcal{E}^d\}_{d=1}^D$ , where  $\mathcal{E}^d$  is the episode for the  $d^{\text{th}}$  patient, and is defined as

$\mathcal{E}^d = (\{y_m^d, t_m^d\}_{m=1}^{M_d}, T_c^d, l^d)$ , with  $y_m^d$  being the  $m^{\text{th}}$   $Q$ -dimensional physiological variable (vital signs and lab test outcomes) for patient  $d$  observed at time  $t_m^d$ , and the total number of samples observed for that patient during her episode is  $M_d$ . The duration of patient  $d$ 's stay in the ward is denoted by  $T_c^d$ , whereas her endpoint outcome (clinical deterioration and ICU admission, or clinical stability and discharge) is declared via a binary variable  $l^d$ ; the realization  $l^d = 1$  means that patient  $d$  was admitted to the ICU, and  $l^d = 0$  means that the patient was discharged home. We stress that the labels  $l^d$  associated with every episode are neither *noisy* nor *entirely subjective*: we assign a label  $l^d = 1$  for patients who actually needed a therapeutic intervention *in* the ICU after an *unplanned* admission, and assign a label  $l^d = 0$  for patients who were discharged and not re-admitted shortly after. We excluded all post-surgical ward patients for whom an ICU admission was preordained since for those patients the prognosis problem is not relevant. Our dataset comprises thousands of episodes for patients admitted to a large medical center over a 3-year period; all the episodes display the structure described above.

### The clinicians were right!

Clinical judgments manifest in the  $d^{\text{th}}$  episode of  $\mathcal{D}$  through informative sampling (encoded in the observation times  $\{t_m^d\}_{m=1}^{M_d}$ ), and informative censoring (encoded in the episode duration  $T_c^d$  and the endpoint outcome  $l^d$ ). Figure 1 is a depiction for both informative sampling and informative censoring. In Figure 1, we estimate the physiological data (time-varying) sampling rate using all the episodes in our dataset over a time horizon of 35 hours before ICU admission for deteriorating patients (i.e. patients with  $l^d = 1$ , with  $t = 0$  being the ICU admission time), and we compute the same estimates for stable patients (i.e. patients with  $l^d = 0$ , with  $t = 0$  being the discharge time).

We can see from the trends in Figure 1 that as the deteriorating patient approaches the ICU admission time, the clinicians tend to sample her physiological data more intensely, whereas as the stable patient approaches the discharge time, the clinicians tend to have a more relaxed schedule for observing her vital signs and lab tests. The divergence between the sampling rates for deteriorating and stable patient groups increases as the patients approach their ICU admission and discharge onsets— that is because the clinicians become less uncertain about the patient's state as time progresses. We have tested the hypothesis that the sampling rate of deteriorating patients is —on average— larger than that for stable patients in the last 24 hours before ICU admission or discharge via a *two-sample t-test* with a significance level of 0.05, and the hypothesis was accepted.

The take-away from Figure 1 is that the clinician's judgment of the patient's clinical state —manifesting in the vital signs and lab tests sampling rates— is very predictive of the




Figure 1. Estimated sampling rates (with 95% confidence intervals) for the physiological data over time.

endpoint outcomes. This implies that there is a room for “learning from the informative clinical judgments”; that is, one can infer the patients’ latent states over time by using the clinicians’ observable sampling patterns as proximal noisy labels for those latent states.

## 3. The Semi-Markov-Modulated Marked Hawkes Process Model

Now we present a probabilistic model for the episodes  $\{\mathcal{E}^d\}_{d=1}^D$  that captures the informative sampling and censoring effects discussed in Section 2. We start by modeling the patient’s latent clinical state process in Subsection 3.1, before modeling the observable variables in Subsection 3.2.

### 3.1. Latent Clinical States

We assume that each patient’s episode is governed by an underlying *latent clinical state process*  $X(t)$  that represents the evolution of her “clinical well-being” over time. The patient’s latent clinical state  $X(t)$  at any point of time  $t \in \mathbb{R}_+$  ( $t = 0$  corresponds to the time of admission to the ward) belongs to a finite space  $\mathcal{X}$  comprising  $N$  states, i.e.  $\mathcal{X} = \{1, 2, \dots, N\}$ . We model informative censoring by assuming that states 1 and  $N$  are *absorbing states*; state 1 is the state of *clinical stability*, at which the patient can be safely discharged home, whereas state  $N$  is the state of *clinical deterioration*, at which the patient needs to be admitted to the ICU. Whenever the patient is in state 1 or  $N$ , her episode is terminated by the clinicians shortly after. All other states in  $\mathcal{X}/\{1, N\}$  are *transient* states in which the patient needs vigilant monitoring by the ward staff.

The clinical state process  $X(t)$  is a *semi-Markov jump process* (Yu, 2010), i.e.  $X(t) = \sum_{n=1}^K X_n \cdot \mathbf{1}_{\{\tau_n \leq t < \tau_{n+1}\}}$  is a semi-Markov process for which every new state realization  $X_n$  starts at a jump time  $\tau_n$ , where  $\tau_1 = 0$ , and lasts for a random sojourn time  $S_n = \tau_{n+1} - \tau_n$ . A total number of  $K$  states are realized in the path  $X(t)$ , where  $K$  is indeed random, and  $X_K \in \{1, N\}$ , i.e. the patient’s episode is concluded by either clinical deterioration or stability. The advantage of adopting a semi-Markovian model for the clinical state process is that unlike Marko-




Figure 2. Depiction of the physiological data sampling times for a clinically deteriorating and a clinically stable patient.

vian models, semi-Markovianity does not imply memoryless transitions—the transition probability from one clinical state to another at any time depends on the time spent in the current state, a property that has been recently validated in various clinical state models (Taghipour et al., 2013). We adopt an *explicit-duration* model for the state sojourn times (Johnson & Willsky, 2013): the  $n^{\text{th}}$  state sojourn time  $S_n$  is drawn from a Gamma distribution<sup>1</sup> with state-specific parameters as follows

$$S_n | (X_n = i) \sim \text{Gamma}(\gamma_i), \forall i \in \mathcal{X}. \quad (1)$$

Since the model includes two absorbing states (1 and  $N$ ) for which the notion of sojourn time is inapplicable, we define the variables  $S_n$  of such states as the clinicians’ “response times” for admitting patients to the ICU or discharging them upon the clinical deterioration/stability onset. The transitions among clinical states are governed by a semi-Markov transition kernel matrix  $\mathbf{P} = (p_{ij})_{i,j}$ , i.e.

$$\mathbb{P}(X_{n+1} = j | X_n = i) = p_{ij}, \quad (2)$$

where self-transitions are eliminated for all transient states (Yu, 2010; Johnson & Willsky, 2013), i.e.  $p_{ii} = 0, \forall i \in \mathcal{X} / \{1, N\}$ , and enforced for the two absorbing states  $p_{ii} = 1, \forall i \in \{1, N\}$ . The initial state distribution is given by  $\pi = (\pi_i)_{i=1}^N$ , where  $\pi_i = \mathbb{P}(X(0) = i)$ . Every episode  $\mathcal{E}^d$  in an EHR dataset  $\mathcal{D}$  is associated with a latent clinical state trajectory  $\{X_n^d, S_n^d\}_{n=1}^{K^d}$ , but we can only observe the absorbing state realization  $X_{K^d} = l^d$  in the EHR data.

### 3.2. Observable Physiological Data

The patient’s latent clinical state process  $X(t)$  manifests in two ways: (1) it modulates the intensity of sampling the patient’s physiological variables (informative sampling), and (2) it modulates the distributional properties of the observed physiological variables. We capture these two effects via a *marked point process* model for the patient’s

<sup>1</sup>We model the sojourn time via a Gamma distribution since it encompasses memoryless exponential distributions of Markov models as a special case (Liu et al., 2015).

episode  $\mathcal{E}$ : the marked point process  $\{(y_m, t_m)\}_{m \in \mathbb{N}_+}$  comprises an *observation process*  $\{t_m\}_{m \in \mathbb{N}_+}$ , which represents the physiological variables’ sampling times, and a *mark process*  $\{y_m\}_{m \in \mathbb{N}_+}$ , which represents the realized physiological variables at these sampling times. The distributional specifications of our marked point process are given in the following Subsections.

#### 3.2.1. THE OBSERVATION PROCESS

We model the observation process generating the physiological variables’ observation times  $\{t_m\}_{m \in \mathbb{N}_+}$  as a *doubly stochastic point process* whose intensity,  $\lambda(t)$ , is a stochastic process modulated by the latent clinical state process  $X(t)$ . In particular, the observation process  $\{t_m\}_{m \in \mathbb{N}_+}$  is modeled as a one-dimensional *Hawkes process* with a linear *self-exciting intensity function*  $\lambda(t, X(t))$  (Lee et al., 2016), i.e.

$$\lambda(t, X(t) = i) = \lambda_i^o + \alpha_i \sum_{\tau < t_m < t} e^{-\beta_i(t-t_m)}, \quad (3)$$

$\forall i \in \mathcal{X}$ , where  $\lambda_i^o$ ,  $\alpha_i$  and  $\beta_i$  are the state-dependent intensity parameters,  $e^{-\beta_i(t-t_m)}$  is an exponential *triggering kernel*, and  $\tau < t$  is the time of the most recent jump in  $X(t)$ . In order to ensure the *local stationarity* of the Hawkes process within the sojourn time of every latent state, we assume that  $\frac{\alpha_i}{\beta_i} < 1, \forall i \in \mathcal{X}$  (Roueff et al., 2016); the expected value of the intensity function is therefore given by  $\mathbb{E}[\lambda(t, X(t) = i)] = \frac{\lambda_i^o}{1 - \frac{\alpha_i}{\beta_i}}$ . For  $\beta_i = \infty$  or  $\alpha_i = 0$ , we recover a modulated *Poisson* process as a special case (Pan et al., 2016).

In Figure 2, we depict the observation times in two patients’ episodes: *patient A* being a clinically deteriorating patient, and *patient B* being a clinically stable patient. We can see that patient A’s episode had its sampling rate escalating as her condition was worsening; the sampling rate remained intense after she was admitted to the ICU. On the other hand, patient B’s episode exhibited a decelerating sampling rate on her path to clinical stability. We can also notice that

the observation times display a subtle “clustered” pattern that point out to their temporal dependencies—indeed, the clinicians are not memoryless, and the times at which they observe the physiological data are dependent. In the light of the above, the modulated Hawkes process described above appears to be a sensible model for the observation process  $\{t_m\}_{m \in \mathbb{N}_+}$  as it captures both the time-varying intensity and the temporal dependencies illustrated in Figure 2.

### 3.2.2. THE MARK PROCESS

Now we provide the distributional specification of the *mark process*  $\{y_m\}_{m \in \mathbb{N}_+}$ . Since the physiological data are irregularly sampled from an underlying continuous-time physiological process  $Y(t)$  at the sampling times determined by the *observation process*  $\{t_m\}_{m \in \mathbb{N}_+}$ , a convenient model for  $Y(t)$  is a *switching Gaussian Process* defined as follows:  $Y(t) = \sum_{n=1}^K Y_n(t) \mathbf{1}_{\{\tau_n \leq t < \tau_{n+1}\}}$ , with

$$Y_n(t) | (X_n = i) \sim \mathcal{GP}(m_i(t), k_i(t, t')), \quad (4)$$

where  $m_i(t)$  and  $k_i(t, t')$  are the state-dependent mean function and covariance kernel, respectively. We use a constant mean function  $m_i(t) = m_i$  and a *Matérn kernel* given by

$$k_i(t, t') = \frac{\left( \frac{\sqrt{2v_i - 1} |t - t'|}{\ell_i} \right)^{v_i - \frac{1}{2}} K_{v_i - \frac{1}{2}} \left( \frac{\sqrt{2v_i - 1} |t - t'|}{\ell_i} \right)}{2^{v_i - \frac{3}{2}} \Gamma(v_i - \frac{1}{2})}, \quad (5)$$

where  $v_i \in \mathbb{N}_+$ ,  $\ell_i \in \mathbb{R}_+$ ,  $\Gamma(\cdot)$  is the *Gamma function* and  $K_{v_i - \frac{1}{2}}(\cdot)$  is a *modified Bessel function* (Rasmussen, 2006).

Our choice for the *Matérn kernel* is motivated by its ability to represent various commonly used stochastic processes; for instance, when  $v_i = 1$ , then  $Y_n(t) | (X_n = i)$  is an *Ornstein-Uhlenbeck process* (Rasmussen, 2006), whereas for a general integer value of  $v_i$ ,  $Y_n(t) | (X_n = i)$  is a continuous-time analogue of the *Auto-regressive process*  $\text{AR}(v_i)$ —a process that has been widely used to model physiological time-series data (Stanculescu et al., 2014). By constructing  $Y_n(t)$  as a continuous-time analog of the AR model, the process  $Y(t) = \sum_{n=1}^K Y_n(t) \mathbf{1}_{\{\tau_n \leq t < \tau_{n+1}\}}$  becomes a continuous-time *switching AR model* that is modulated by the patient’s latent clinical state process  $X(t)$ . We observe the continuous-time process  $Y(t)$  only at the sampling times dictated by the observation process  $\{t_m\}_{m \in \mathbb{N}_+}$ , and the resulting process  $\{y_m\}_{m \in \mathbb{N}_+}$  defines the mark process. The observation process together with the mark process, both modulated by the latent clinical state process  $X(t)$ , constitute a *marked Hawkes process*, which completely describes a patient’s episode.

The mark process defined above is one-dimensional, and hence we need to extend the definition to handle a multi-dimensional process that represents multiple lab tests and

vital signs. In other words, we seek a continuous-time analog of the *switching Vector Auto-regressive (VAR) model* rather than an AR model<sup>2</sup>. This is achieved by adopting a *multi-task Gaussian Process* as a model for the  $Q$ -dimensional physiological process  $Y(t) \in \mathbb{R}^Q$  (Durichen et al., 2015). That is, we assume that  $Y_n(t) | (X_n = i) \sim \mathcal{GP}(m_i(t), \mathbf{K}_i(t, t'))$ , where the covariance kernel  $\mathbf{K}_i(t, t') = \{k_i(r, g, t, t')\}_{r, g=1}^Q$  is based on the *intrinsic correlation model* (Bonilla et al., 2007), i.e.  $\mathbf{K}_i(t, t')$  can be written in the following separable form

$$k_i(r, g, t, t') = \Sigma_i(r, g) \cdot k_i(t, t'), \quad (6)$$

where  $k_i(r, g, t, t')$  is the covariance between the  $r^{\text{th}}$  physiological variable at time  $t$  and the  $g^{\text{th}}$  physiological variable at time  $t'$ ,  $\Sigma_i$  is a  $Q \times Q$  *intrinsic correlation matrix*, and  $k_i(t, t')$  is the Matérn kernel in (5). For every state  $i$ , we denote the multi-task GP parameter set as  $\Theta_i = (m_i(t), \mathbf{K}_i(t, t'))$ , and the Hawkes process parameter set as  $\Lambda_i$ . The entire model parameters can be bundled in the parameter set  $\Omega$  as follows

$$\Omega = \{\mathbf{P}, (\pi_i, \gamma_i, \Lambda_i, \Theta_i)_{i \in \mathcal{X}}\}.$$

Given a parameter set  $\Omega$ , we can easily generate sample patient episodes from our model by first sampling a state sequence  $\{X_1, \dots, X_K\}$  using the semi-Markov transition kernel, then sampling a corresponding sequence of sojourn times  $\{S_1, \dots, S_K\}$  from the state-dependent Gamma distributions, then sampling a set of multi-task Gaussian process  $\{Y_1(t), \dots, Y_K(t)\}$ , and finally sampling a sequence of observation times  $\{t_m\}_{m=1}^M$  using Ogata’s *modified thinning algorithm* (Ogata, 1981). Figure 3 depicts one patient episode sampled from our model. (An algorithm for sampling episodes from our model is provided in Appendix B in the supplementary material.)

## 4. Learning and Inference

In this Section, we develop an *offline learning algorithm* that learns the model parameter  $\Omega$  using the offline training episodes in  $\mathcal{D}$ , and a *real-time risk scoring algorithm* that computes a hospitalized patient’s risk over time. The learning algorithm operates by first detecting change-points in the physiological data and the observation process; using the detected change-points, the algorithm segments each episode into a sequence of states, and uses an EM algorithm to estimate the model parameters. The real-time risk scoring algorithm operates by inferring the patient’s current state via forward-filtering, and then computing the probability of eventual absorption in the deteriorating state.

<sup>2</sup>In Appendix A of the supplementary material, we establish the connection between the switching multi-task Gaussian Process model described herein and the conventional VAR model, showing that the former is the continuous-time analog of the latter.




Figure 3. An episode sampled from the proposed model.

#### 4.1. The Offline Learning Algorithm

Estimating the model parameters  $\Omega$  from the dataset  $\mathcal{D} = \{\mathcal{E}^d\}_{d=1}^D$  is a daunting task due to the hiddenness of the patients' clinical state trajectories  $\{X_n^d, S_n^d\}_{d=1}^D$ ; an application of MCMC-based inference methods, such as the method in (Qin & Shelton, 2015), will incur an excessive computational cost for a complex model like ours. We therefore developed an efficient three-step learning algorithm that capitalizes on the structure of the patients' episodes in order to find a point estimate  $\hat{\Omega}$  for  $\Omega$ . The three steps of the our learning algorithm are listed hereunder<sup>3</sup>.

##### Step 1: Change-point detection

We first estimate the jump times  $\{\tau_1^d, \dots, \tau_{K^d}^d\}$  for every episode  $\mathcal{E}^d$  in  $\mathcal{D}$ . This is achieved by using the *E-divisive* change-point detection algorithm (Matteson & James, 2014). Since the *E-divisive* algorithm is nonparametric, we are able to estimate the onsets of all clinical states, i.e. the jump times of  $X^d(t)$ , prior to finding the estimate  $\hat{\Omega}$ . We let the *E-divisive* algorithm jointly detect changes in the distributions of the observable variables  $\{y_m^d\}_{m=1}^{M^d}$  and the observation process  $\{t_m^d\}_{m=1}^{M^d}$  by creating an *augmented* vector of observables that comprises both the physiological observations and a “differential” observation process, i.e.

$$\{\hat{\tau}_1^d, \dots, \hat{\tau}_{K^d}^d\} = \text{E-divisive}((y_1^d, \Delta t_1^d), \dots, (y_{M^d}^d, \Delta t_{M^d}^d)),$$

where  $\Delta t_m^d = t_m^d - t_{m-1}^d$ , with  $\Delta t_1^d = 0$ . By the end of this step, we obtain an estimate for the start and end times of all clinical state realizations for every episode  $\mathcal{E}^d$ .

##### Step 2: Maximum Likelihood Estimation (MLE) of the absorbing states' parameters

By virtue of informative censoring, we know the identities of all the absorbing states, i.e.  $X_{K^d}^d = l^d, \forall d$ . Now that we have estimates for the onsets of the absorbing states, obtained from step 1, then we can estimate the response

<sup>3</sup>The details for all the algorithms in this Section are provided in Appendix C in the supplementary material.

times as  $\hat{S}_{K^d}^d = T_c^d - \hat{\tau}_{K^d}^d, \forall 1 \leq d \leq D$ . Define the (fully observable) sub-dataset  $\mathcal{D}_i$  as follows

$$\mathcal{D}_i = \left\{ (y_m, t_m)_{\{t_m \geq \hat{\tau}_{K^d}^d\}}, \hat{S}_{K^d}^d : l^d = i, \mathcal{E}^d \in \mathcal{D} \right\},$$

$\forall i \in \{0, 1\}$ . Given such a fully fledged specification of the absorbing states' onsets, identities, and the corresponding physiological variables, we can directly apply MLE to estimate the parameters  $\Theta_1, \Theta_N, \gamma_1, \gamma_N, \Lambda_1$  and  $\Lambda_N$ . Using the dataset  $\mathcal{D}_i$ , the parameter  $\Theta_i$  is estimated using the gradient method for Gaussian processes as in (Bonilla et al., 2007),  $\gamma_i$  is estimated using the standard MLE estimating equations, and  $\Lambda_i$  is estimated by maximizing the recursive likelihood formula in (Ogata, 1981) using the Nelder-Mead simplex method (Nelder & Mead, 1965).

##### Step 3: Estimation of the transient states' parameters using the EM algorithm

While the absorbing states are observable, the transient states are all hidden. In order to estimate the parameters  $\mathbf{P}, \{\Theta_i\}_{i=2}^{N-1}, \{\gamma_i\}_{i=2}^{N-1}$ , and  $\{\Lambda_i\}_{i=2}^{N-1}$ , we use the jump times' estimates  $\{\hat{\tau}_1^d, \dots, \hat{\tau}_{K^d-1}^d\}$  (obtained from step 1) in order to *segment* every episode  $d$  into a set of finite transition, and hence we obtain a discrete-time HMM-like process. We truncate all the episodes by removing the data belonging to the absorbing state, and run the EM algorithm in order to estimate the transient state parameters. The EM algorithm takes advantage of informative censoring in the forward-backward message passing stage by computing the backward messages conditioned on the identity of the endpoint absorbing state  $l^d$  for every episode  $d$ .

#### 4.2. The Real-time Risk Scoring Algorithm

Having learned the model parameters  $\hat{\Omega}$  from an offline dataset  $\mathcal{D}$ , we now explain how risk scoring is conducted in real-time for a newly hospitalized patient. The patient risk score at time  $t$  is denoted by  $R(t)$ , and is defined as  $R(t) = \mathbb{P}(X(\infty) = N | \{y_m, t_m\}, t_m \leq t, \hat{\Omega})$ . That is, the risk score  $R(t)$  is the probability of being eventually absorbed in the deteriorating state  $N$  given the observable physiological data up to time  $t$ . Using Bayes' rule, the risk score  $R(t)$  is given by

$$\sum_{i \in \mathcal{X}} \underbrace{\mathbb{P}(X(t) = i | \{y_m, t_m\})}_{\text{Current state}} \cdot \underbrace{\mathbb{P}(X(\infty) = N | X(t) = i)}_{\text{Future transition}},$$

where the “current state” term is computed in real-time using the efficient (dynamic programming) *forward-filtering* algorithm, whereas the “future transition” term is computed offline using the estimated model parameters.

## 5. Experiments

In order to evaluate the prognostic utility of our model, we conducted experiments on a dataset  $\mathcal{D}$  comprising informa-

tion on patient admissions to a major medical center over a 3-year period, and compared the proposed risk score defined in Subsection 4.2 with other competing baselines. We briefly describe our dataset in the next Subsection, and then present the experimental results. A very detailed description for our dataset and the implementation of the baselines is provided in Appendix D in the supplementary material.

### 5.1. Data Description

Each patient record in  $\mathcal{D}$  is an episode that is formatted as described in Section 2. The patients' cohort in  $\mathcal{D}$  is very heterogeneous; diagnoses included sepsis, hypertension, renal failure, leukemia, septicemia and pneumonia. Each patient's episode in  $\mathcal{D}$  comprises 21 vital signs and lab tests that are collected for the patient over time, along with the time instances at which they were collected. The vital signs include diastolic and systolic blood pressure, Glasgow coma scale score, heart rate, eye opening, respiratory rate, temperature,  $O_2$  saturation and device assistance, best motor and verbal responses. The lab tests included measurements of chloride, glucose, urea nitrogen, white blood cell count, creatinine, hemoglobin, platelet count, potassium, sodium and  $CO_2$ . In all the experiments conducted in this Section, we split  $\mathcal{D}$  into a training set with admissions over a 2.5-year period (5,000 episodes) and a testing set with admissions over a 6-month period (1,094 episodes).

### 5.2. Results

We ran the offline learning algorithm in Subsection 4.1 (with 1000 EM iterations) on the 5,000 training episodes in  $\mathcal{D}$ , and obtained an estimate  $\hat{\Omega}$  for the semi-Markov-modulated marked Hawkes process that describes the patient cohort. Using the *Bayesian information criterion*, we selected an instantiation of our model with 4 clinical states, where state 1 is the *clinical stability* (absorbing) state, and state 4 is the *clinical deterioration* (absorbing) state. The learned Hawkes process intensity functions for these two states are given by

$$\lambda(t, 1) = \underbrace{0.55}_{\text{Baseline intensity } \downarrow} + \underbrace{0.2 \sum_{t_m} e^{-8.46(t-t_m)}}_{\text{Temporal dependencies } \downarrow},$$

$$\lambda(t, 4) = \underbrace{0.82}_{\text{Baseline intensity } \uparrow} + \underbrace{0.16 \sum_{t_m} e^{-1.36(t-t_m)}}_{\text{Temporal dependencies } \uparrow},$$

where  $\lambda(t, X(t))$  is measured in *samples per hour*. We note that the estimated Hawkes process parameters accurately describe the clinicians' judgments; when the patient is in the deteriorating state (state 4), the clinicians tend to observe her physiological measurements more frequently. This manifests in the baseline intensity of  $\lambda(t, 4)$  being




Figure 4. An episode for a hospitalized cardiac patient.

50% higher than that of  $\lambda(t, 1)$ . Moreover, we note that when the patient is clinically stable (state 1), the temporal dependencies between the observation times almost disappear as the exponential triggering kernel plays little role in determining the observation times, i.e.  $\lambda(t, 1) \approx 0.55$ , which renders the observation process closer to a Poisson process. Contrarily, when the patient is deteriorating, strong temporal dependencies are displayed in the observation times—this is intuitive since for a deteriorating patient, the follow-up times decided by the clinicians strongly depend on what have been observed in the past. These distinguishing state-specific features of the clinical judgments are the essence of informative sampling, which allows us to integrate physiological data together with clinical experience while learning the patient's physiological model.

We illustrate the value of informative sampling in Figure 4 through an episode for a cardiac patient who was admitted in the ward for 1 week before being sent to the ICU upon a cardiac arrest. When running the offline learning algorithm in Section 4.1 while assuming that the observation process  $\{t_m^d\}_m$  is uninformative (i.e.  $\lambda(t, i) = \lambda(t, j), \forall i, j \in \mathcal{X}$ ), the detected clinical deterioration onset is only 10 hours ahead of the cardiac arrest event (onset (2) in Figure 4): this is because the states are estimated solely based on the physiological data, and hence the clinical deterioration state is detected only when the patient's heart rate fell below 60 beats per minute (BPM). When running the learning again but with informative sampling taken into account, the detected clinical deterioration onset is 40 hours ahead of the cardiac arrest event (onset (1) in Figure 4); this is the time instance at which the clinicians decided to monitor the patient more vigilantly (i.e. more intense sampling rate) even though the evidence for an upcoming cardiac arrest in her heart rate trajectory was rather subtle. The clinicians' decision to intensely observe the patient's physiological trajectory is based either on their experience, or on apparent symptoms or complains from the patient that were not recorded in the EHR. By integrating the clinicians' judgments into our model, we are able to capture the subtleties in the patients' temporal physiological parameters, and hence learn more accurate representations for the clinical states.

Method		AUROC
Proposed Score		0.481
Medical Scores	Rothman Index	0.255
	MEWS	0.182
	APACHE-II	0.130
	SOFA	0.127
ML Algorithms	SW-RF	0.366
	HMMs	0.321
	SW-GP	0.305
	SW-LR	0.267

Table 1. Comparison between various risk scoring methods ( $p < 0.01$ ). (SW: Sliding-window, GP: Gaussian process regression, LR: Logistic Regression, RF: Random Forests)

We then computed the real-time risk score (as described in Subsection 4.2) for the testing episodes, and compared its *sensitivity-precision* AUCROC with that of the baseline risk scoring methods listed hereunder.

**Medical Risk Scores:** we considered the two most commonly used medical risk scores in regular wards– the MEWS score (Morgan et al., 1997) and the Rothman index<sup>4</sup> (Rothman et al., 2013). We implemented the MEWS score and the Rothman index as described in (Kyriacos et al., 2014) and (Rothman et al., 2013), respectively. We also compared with the APACHE-II and SOFA scores.

**Machine Learning Algorithms:** we considered the traditional approach for real-time risk scoring, which treats informative censoring as a surrogate label based on which a supervised regression model is learned offline, and then risk scoring is applied in real-time using the temporal data within a sliding window– we call these methods “sliding-window methods”. We implemented sliding-window methods based on logistic regression (Ho et al., 2012; Saria et al., 2010), random forests, and Gaussian process regression (Ghassemi et al., 2015; Yoon et al., 2016). In addition, we compared our risk score with a standard Hidden Markov Model with Gaussian emissions. The relevant physiological measurements for every baseline were selected through the *correlated feature selection* method (Yu & Liu, 2003). The hyper-parameters of all the baselines, including the size of the sliding window for the supervised learning methods, were optimized via cross-validation. To handle the irregularly sampled data, we discretized the time horizon into 1-hour steps and fed the baselines with interpolated, discrete-time episodes.

In Table 1, we compare the performance of our risk score with the baselines in terms of the sensitivity-precision AU-

<sup>4</sup>At the time of conducting these experiments, the Rothman index was deployed in more than 60 major hospitals in the US.

ROC. As we can see, the proposed risk score offers a 23% AUROC improvement as compared to the best performing medical risk score –the Rothman index– which was the score deployed in our medical center at the time of conducting this experiment. Moreover, our risk score also provides significant gains over discriminative sliding-window regression models; the proposed risk score achieves a 11.5% AUROC improvement as compared to the best performing ML algorithm (random forest). In addition, the proposed risk score achieves a 16% AUROC improvement as compared to a standard HMM with Gaussian emission variables<sup>5</sup>. On average, our risk score prompts ICU alarms 8 hours before the censoring time at a sensitivity of 50% and precision of 35%.

## 6. Discussion: Chicken-and-egg

We stress that while computing our risk score for the testing episodes, we did not use the information conveyed in the observation process  $\{t_m^d\}_m$  to infer the patients’ clinical states. This is because in practice, the value of the real-time risk score  $R(t)$  itself influences the clinician’s behavior and hence impacts the observation process, creating a *chicken-and-egg* dilemma in which one cannot clearly conceptualize the causal relation between the risk score and the clinicians’ judgments. A very interesting research direction is to consider an observation process that is modulated by both the patient’s state and the real-time risk score through an intensity function  $\lambda(t, X(t), R(t))$ , where the algorithm learns the clinical state representation *online* by “sharing experience” with the clinicians. That is, the algorithm uses the clinician’s judgments to refine its clinical state model, which leads to a refined risk score  $R(t)$  that would in turn allow the clinician to exhibit more accurate judgments; an online learning process that would ideally converge to a state of “shared knowledge” between the clinician and the system.

The significant prognostic value offered by our risk score promises a great improvement in the quality of subacute care in wards. By utilizing the proposed score instead of the current technology, clinicians in a crowded ward can better focus their attention on patients at real risk of deterioration, and can also plan for timely ICU admissions and effective therapeutic interventions. With the high in-hospital mortality rates in wards, deploying our risk score may help save thousands of lives annually– we are currently working towards installing the proposed risk score in our medical center.

<sup>5</sup>The adoption of a semi-Markovian model for the clinical state process protects our model from the overtly rapid state switching behavior that is introduced by memoryless HMMs (Matteson & James, 2014).



## References

- Bonilla, Edwin V, Chai, Kian M, and Williams, Christopher. Multi-task gaussian process prediction. In *Advances in neural information processing systems*, pp. 153–160, 2007.
- Charles, D, Gabriel, M, and Henry, J. Electronic capabilities for patients among us non-federal acute care hospitals: 2008-2014. onc data brief 29. office of the national coordinator for health information technology: Washington dc, 2016.
- Che, Zhengping, Purushotham, Sanjay, Cho, Kyunghyun, Sontag, David, and Liu, Yan. Recurrent neural networks for multivariate time series with missing values. *arXiv preprint arXiv:1606.01865*, 2016.
- Clifton, Lei, Clifton, David A, Pimentel, Marco AF, Watkinson, Peter J, and Tarassenko, Lionel. Gaussian process regression in vital-sign early warning systems. In *Engineering in Medicine and Biology Society (EMBC), 2012 Annual International Conference of the IEEE*, pp. 6161–6164. IEEE, 2012.
- Cvach, Maria. Monitor alarm fatigue: an integrative review. *Biomedical instrumentation and technology*, 46(4):268–277, 2012.
- Du, Nan, Dai, Hanjun, Trivedi, Rakshit, Upadhyay, Utkarsh, Gomez-Rodriguez, Manuel, and Song, Le. Recurrent marked temporal point processes: Embedding event history to vector. In *Proceedings of the 22nd ACM SIGKDD*, pp. 1555–1564. ACM, 2016.
- Durichen, Robert, Pimentel, Marco AF, Clifton, Lei, Schweikard, Achim, and Clifton, David A. Multitask gaussian processes for multivariate physiological time-series analysis. *Biomedical Engineering, IEEE Transactions on*, 62(1):314–322, 2015.
- Finlay, G Duncan, Rothman, Michael J, and Smith, Robert A. 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*, 9(2):116–119, 2014.
- Ghassemi, Marzyeh, Pimentel, Marco AF, Naumann, Tris-tan, Brennan, Thomas, Clifton, David A, Szolovits, Peter, and Feng, Mengling. A multivariate timeseries modeling approach to severity of illness assessment and forecasting in icu with sparse, heterogeneous clinical data. In *AAAI*, pp. 446–453, 2015.
- Hawkes, Alan G and Oakes, David. A cluster process representation of a self-exciting process. *Journal of Applied Probability*, pp. 493–503, 1974.
- Henry, Katharine E, Hager, David N, Pronovost, Peter J, and Saria, Suchi. A targeted real-time early warning score (trewscore) for septic shock. *Science Translational Medicine*, 7(299):299ra122–299ra122, 2015.
- Ho, Joyce C, Lee, Cheng H, and Ghosh, Joydeeph. Imputation-enhanced prediction of septic shock in icu patients. In *Proceedings of the ACM SIGKDD Workshop on Health Informatics*, pp. 21–27, 2012.
- Hogan, Helen, Healey, Frances, Neale, Graham, Thomson, Richard, Vincent, Charles, and Black, Nick. Preventable deaths due to problems in care in english acute hospitals: a retrospective case record review study. *BMJ quality & safety*, pp. bmjqs–2012, 2012.
- Johnson, Alistair EW, et. al. Mimic-iii, a freely accessible critical care database. *Scientific data*, 3, 2016.
- Johnson, Matthew J and Willsky, Alan S. Bayesian non-parametric hidden semi-markov models. *Journal of Machine Learning Research*, 14(Feb):673–701, 2013.
- Kause, Juliane, Smith, Gary, Prytherch, David, Parr, Michael, Flabouris, Arthas, Hillman, Ken, et al. A comparison of antecedents to cardiac arrests, deaths and emergency intensive care admissions in australia and new zealand, and the united kingdomthe academia study. *Resuscitation*, 62(3):275–282, 2004.
- Kumar, Gagan, Kumar, Nilay, Taneja, Amit, Kaleekal, Thomas, Tarima, Sergey, McGinley, Emily, Jimenez, Edgar, Mohan, Anand, Khan, Rumi Ahmed, Whittle, Jeff, et al. Nationwide trends of severe sepsis in the 21st century (2000-2007). *Chest Journal*, 140(5):1223–1231, 2011.
- Kyriacos, Una, Jelsma, Jennifer, James, Michael, and Jordan, Sue. Monitoring vital signs: development of a modified early warning scoring (mews) system for general wards in a developing country. *PloS one*, 9(1):e87073, 2014.
- Lee, Young, Lim, Kar Wai, and Ong, Cheng Soon. Hawkes processes with stochastic excitations. pp. 79–88, 2016.
- Liu, Yu-Ying, Li, Shuang, Li, Fuxin, Song, Le, and Rehg, James M. Efficient learning of continuous-time hidden markov models for disease progression. In *Advances in neural information processing systems*, pp. 3600–3608, 2015.
- Mardini L, Lipes J, Jayaraman D. Adverse outcomes associated with delayed intensive care consultation in medical and surgical inpatients. *Journal of critical care*, 27(6):688–693, 2012.

- Matteson, David S and James, Nicholas A. A nonparametric approach for multiple change point analysis of multivariate data. *Journal of the American Statistical Association*, 109(505):334–345, 2014.
- Merchant, Raina M, Yang, Lin, Becker, Lance B, Berg, Robert A, Nadkarni, Vinay, Nichol, Graham, Carr, Brendan G, Mitra, Nandita, Bradley, Steven M, Abella, Benjamin S, et al. Incidence of treated cardiac arrest in hospitalized patients in the united states. *Critical care medicine*, 39(11):2401, 2011.
- Mokart, Djamel et. al. Delayed intensive care unit admission is associated with increased mortality in patients with cancer with acute respiratory failure. *Leukemia & lymphoma*, 54(8):1724–1729, 2013.
- Morgan, RJM, Williams, F, and Wright, MM. An early warning scoring system for detecting developing critical illness. *Clin Intensive Care*, 8(2):100, 1997.
- Moskovitch, Robert, Walsh, Colin, Wang, Fei, Hripcsak, George, and Tatonetti, Nicholas. Outcomes prediction via time intervals related patterns. In *Data Mining (ICDM), 2015 IEEE International Conference on*, pp. 919–924. IEEE, 2015.
- Nelder, John A and Mead, Roger. A simplex method for function minimization. *The computer journal*, 7(4):308–313, 1965.
- Nguyen, H Bryant, Corbett, Stephen W, Steele, Robert, Banta, Jim, Clark, Robin T, Hayes, Sean R, Edwards, Jeremy, Cho, Thomas W, and Wittlake, William A. Implementation of a bundle of quality indicators for the early management of severe sepsis and septic shock is associated with decreased mortality. *Critical care medicine*, 35(4):1105–1112, 2007.
- Ogata, Yoshihiko. On lewis’ simulation method for point processes. *IEEE Transactions on Information Theory*, 27(1):23–31, 1981.
- Pan, Jiangwei, Rao, Vinayak, Agarwal, Pankaj K., and Gelfand, Alan E. Markov-modulated marked poisson processes for check-in data. In *Proceedings of ICML*, pp. 2244–2253, 2016.
- Parshuram, Christopher S, Hutchison, James, and Midgough, Kristen. Development and initial validation of the bedside paediatric early warning system score. *Crit Care*, 13(4):R135, 2009.
- Qin, Zhen and Shelton, Christian R. Auxiliary gibbs sampling for inference in piecewise-constant conditional intensity models. In *UAI*, pp. 722–731, 2015.
- Rasmussen, Carl Edward. Gaussian processes for machine learning. 2006.
- Rothman, Michael J, Rothman, Steven I, and Beals, Joseph. Development and validation of a continuous measure of patient condition using the electronic medical record. *Journal of biomedical informatics*, 46(5): 837–848, 2013.
- Roueff, François, Von Sachs, Rainer, and Sansonnet, Laure. Locally stationary hawkes processes. *Stochastic Processes and their Applications*, 126(6):1710–1743, 2016.
- Saria, Suchi, Rajani, Anand K, Gould, Jeffrey, Koller, Daphne, and Penn, Anna A. Integration of early physiological responses predicts later illness severity in preterm infants. *Science translational medicine*, 2(48):48ra65–48ra65, 2010.
- Stanculescu, Ioan, Williams, Christopher KI, and Freer, Yvonne. Autoregressive hidden markov models for the early detection of neonatal sepsis. *IEEE journal of biomedical and health informatics*, 18(5):1560–1570, 2014.
- Taghipour, S, Banjevic, D, Miller, AB, Montgomery, N, Jardine, AKS, and Harvey, BJ. Parameter estimates for invasive breast cancer progression in the canadian national breast screening study. *British journal of cancer*, 108(3):542–548, 2013.
- Wang, Xiang, Sontag, David, and Wang, Fei. Unsupervised learning of disease progression models. In *Proceedings of the 20th ACM SIGKDD*, pp. 85–94. ACM, 2014.
- Wiens, Jenna, Horvitz, Eric, and Guttag, John V. Patient risk stratification for hospital-associated c. diff as a time-series classification task. In *Advances in Neural Information Processing Systems*, pp. 467–475, 2012.
- Yoon, J, Alaa, A, Hu, S, and van der Schaar, M. Forecasticu: A prognostic decision support system for timely prediction of intensive care unit admission. *Proceedings of ICML*, pp. 1680–1689, 2016.
- Yu, Lei and Liu, Huan. Feature selection for high-dimensional data: A fast correlation-based filter solution. In *ICML*, volume 3, pp. 856–863, 2003.
- Yu, Shun, Leung, Sharon, Heo, Moonseong, Soto, Graciela J, Shah, Ronak T, Gunda, Sampath, and Gong, Michelle Ng. Comparison of risk prediction scoring systems for ward patients: a retrospective nested case-control study. *Critical Care*, 18(3):1, 2014.
- Yu, Shun-Zheng. Hidden semi-markov models. *Artificial Intelligence*, 174(2):215–243, 2010.