
ForecastICU: A Prognostic Decision Support System for Timely Prediction of Intensive Care Unit Admission

Jinsung Yoon¹
Ahmed M. Alaa¹
Scott Hu², MD
Mihaela van der Schaar¹

JSYOON0823@UCLA.EDU
AHMEDMALAA@UCLA.EDU
SCOTTHU@MEDNET.UCLA.EDU
MIHAELA@EE.UCLA.EDU

¹Electrical Engineering Department, University of California, Los Angeles (UCLA), CA, 90095, USA.

²David Geffen School of Medicine, University of California, Los Angeles (UCLA), CA, 90095, USA.

Abstract

We develop *ForecastICU*: a prognostic decision support system that monitors hospitalized patients and prompts alarms for intensive care unit (ICU) admissions. ForecastICU is first trained in an offline stage by constructing a Bayesian *belief system* that corresponds to its belief about how trajectories of physiological data streams of the patient map to a clinical status. After that, ForecastICU monitors a new patient in real-time by observing her physiological data stream, updating its *belief* about her status over time, and prompting an alarm whenever its *belief process* hits a predefined threshold (confidence). Using a real-world dataset obtained from UCLA Ronald Reagan Medical Center, we show that ForecastICU can predict ICU admissions 9 hours before a physician's decision (for a sensitivity of 40% and a precision of 50%). Also, ForecastICU performs consistently better than other state-of-the-art machine learning algorithms in terms of sensitivity, precision, and timeliness: it can predict ICU admissions 3 hours earlier, and offers a 7.8% gain in sensitivity and a 5.1% gain in precision compared to the best state-of-the-art algorithm. Moreover, ForecastICU offers an area under curve (AUC) gain of 22.3% compared to the Rothman index, which is the currently deployed technology in most hospital wards.

1. Introduction

Intensive Care Unit (ICU) admission for hospitalized patients is a vital and delay-critical decision. The timing of transfer to the ICU is an important determinant of a patient's outcomes. Various medical studies have demonstrated that delayed identification of clinical deterioration, leading to delayed admission to the ICU and delayed therapeutic intervention, results in increased morbidity and mortality (Liu et al., 2012) (Young et al., 2003). This illustrates the urgency to develop automated prognostic decision support systems that alert the medical staff of impending clinical deterioration, enabling clinicians to intervene at an earlier time, thereby preventing an arrest or reducing the need for ICU transfer.

An automated system for ICU admission prediction is envisioned to operate in the following manner. The system is fed with high-dimensional physiological data streams that belong to a monitored patient and tries to infer whether the patient is clinically deteriorating or not in a timely manner, i.e. earlier than the time at which a physician would normally decide to impend an ICU transfer for that patient. The system can take advantage of the available electronic health record (EHR) data to learn the trends in the physiological data streams associated with patients who previously got discharged or admitted to the ICU (Churpek et al., 2014).

Designing a system that carries out the steps described above is associated with many practical and technical challenges. First, while data streams for previously hospitalized patients are recorded in the EHR, the *clinical status* of such patients upon their hospitalization *differs* from one patient to another; thus, learning from the labeled patient examples is not straightforward as the examples themselves entail some ambiguity. Second, *not all* the data streams are *relevant* to the ICU admission decision, and not all of the different streams are sampled at the same rate, i.e. some

of the data samples are missing from some streams. Finally, the patients are monitored only for a finite amount of time, and the time-series observed by the system would eventually stop; thus the system should issue the prediction within a reasonable amount of time ahead of an *unknown deadline*, i.e. the actual time when physicians decide to admit the patient to ICU or discharge the patient.

In this paper, we develop ForecastICU, a prognostic decision support system that carries out timely predictions of ICU admissions for hospitalized patients. ForecastICU adopts a Bayesian approach for issuing predictions; by applying density estimation using the data streams of previously hospitalized patients, ForecastICU constructs a *belief system* that corresponds to its belief about the patient’s clinical status as a function of time. ForecastICU approaches a new patient by observing her physiological data stream, updates its *belief* about the patient over time as it observes more samples from her physiological stream, and prompts an ICU admission alarm whenever the system’s *belief process* hits a predefined threshold that quantifies the system’s confidence in the issued alarms.

Our approach hinges on the idea that predicting ICU admissions from temporal physiological streams can be viewed as an optimal stopping problem; or in other words, it is equivalent to learning an *unknown stopping rule* of a stochastic process, i.e. learning how physicians make ICU admission decisions from the recorded temporal data streams (realizations of a stochastic process) of previously hospitalized patients, and hence promptly issuing ICU admission predictions (alarms) before the stopping times of these processes (i.e. before the ICU admission decision that physicians would actually make without ForecastICU). Technically, ForecastICU approaches the ICU prognosis problem as an optimal stopping problem with uncertainty in the initial clinical status and the distribution of the physiological data streams.

We have applied ForecastICU to a real-world dataset obtained from UCLA Ronald Reagan Medical Center. Experiments show that, for a sensitivity of 40% and a precision of 50%, ForecastICU can predict ICU admissions as early as 9 hours (on average) before the actual physician’s decision. Moreover, we show that ForecastICU performs consistently better than other state-of-the-art algorithms for all ranges of the system’s parameters in terms of sensitivity, precision, and timeliness: it can predict ICU admissions 3 hours earlier, and offers a 7.8% gain for sensitivity and 5.1% gain for precision with an 8.15% gain for the area under curve (AUC) compared to the best benchmark algorithm. Such gains can map to significant reductions in ICU mortality rates and better resource utilization in hospitals. Most remarkably, ForecastICU offers significant gains with respect to state-of-the-art risk assessment technologies: if

achieves a 22.3% AUC gain compared to the Rothman index, the currently deployed technology in most hospital wards. This gain translates to 3 times less false alarms issued by ForecastICU as compared to those issued by the Rothman index at a sensitivity of 50%; hence, ForecastICU promises significant improvements in alarm credibility and clinical resource management if deployed in hospital wards in replacement of current technologies.

2. Related works

2.1. ICU prognosis

Methods for supporting prognostic clinical decisions have been investigated both in the medical literature and in the data mining literature. Several clinical studies have investigated the effectiveness of the usage of early warning scores (EWS) for the detection of patient deterioration (Bruijns et al., 2013). Such methods identify hospitalized patients who reach a certain “trigger” threshold (Churpek et al., 2014) (Prytherch et al., 2010), and consequently alarms the clinicians to impend an ICU transfer for those patients. However, since EWS-based methods respond to triggering events that may not signal a truly deteriorating patient, they suffer from high rates of false alarms (70-95%) (Tsien & Fackler, 1997), which results in alarm fatigue and inappropriate resource utilization. Recent systematic reviews have demonstrated that EWS-based alarms only marginally improve outcomes while substantially increasing physician and nursing workloads (Alam et al., 2014).

Prognostic decision support has been investigated in the data mining literature as well. In (Steyerberg et al., 2001), (Seker et al., 2003), and (Steyerberg et al., 2000), simple regression models were developed to carry out risk assessment for developing diseases like breast and prostate cancer based on the patients’ features. However, such predictive models deal with scenarios where predictions span years rather than hours, i.e. the delay-sensitivity of such models is much coarser and less critical. Prognostic decision support for delay-critical settings has been addressed in the context of ICU environments in (Sun et al., 2010). (Neti et al., 2010) and (Zhang et al., 2012). However, the objective in these models was to predict the trajectory of physiological data streams using previously hospitalized patients’ synchronized data streams. Unlike ForecastICU, the models therein have developed simple auto-regressive models to extrapolate incoming data streams, but have not addressed the problem of timely classification of incoming patients whose monitored data streams entail an unknown stopping time, and are not synchronized with the training data.

2.2. Time-series Analysis

Predicting ICU admissions based on physiological data streams is essentially a time-series analysis problem. In this subsection, we position our work with respect to two different strands of the time-series analysis literature that are relevant to the ICU prediction problem; namely, time-series classification and temporal physiological stream modeling.

Time-series classification: Most of the state-of-the-art time-series classification algorithms rely on clustering the time-series in the training data using a distance metric and then associating the new incoming series with one of the learned clusters (Xing et al., 2012; Wei & Keogh, 2006; Buza et al., 2010; Hüsken & Stagge, 2003; Van Heeswijk et al., 2009). Our work departs from this literature in the following aspects. First, unlike the conventional time-series classification problem, the physiological data streams in our problem stop at some (random) point of time (when the patient is discharged or admitted to ICU). This requires both introducing new design objectives, i.e. timeliness, and accounting for the stochastic deadline for the classification task (i.e. clinical deterioration and the clinicians’ ICU admission decisions). Second, our training data comprises data streams for patients who were hospitalized with different clinical statuses; this requires learning how to “align” the training data to construct a functioning time-series model. Both issues were not addressed by the previous time-series classification algorithms, which operate on synchronized and fixed-length streams.

Temporal physiological stream modeling: Previous attempts have been made to model physiological data streams, mainly using Gaussian Processes (GP) in (Clifton et al., 2012), (Ghassemi et al., 2015)(Durichen et al., 2015), and (Pimentel et al., 2013). In all these works, the focus is to predict the values of the physiological data via GP regression (e.g. estimating cerebrovascular pressure reactivity in (Ghassemi et al., 2015)), but not to make *timely* decisions. Moreover, such models have been limited to the usage of the *squared-exponential* covariance kernel (e.g. see eq. (2) in (Ghassemi et al., 2015)), which can only captures stationary physiological; for modeling the streams of clinically deteriorating patients, i.e. patients who need ICU admission, we need to consider non-stationary models as well.

3. Theoretical foundations for ForecastICU

In this section, we provide a formal model for an abstract *forecaster* that aims at learning to predict ICU admissions. The theoretical analysis of such a forecaster will serve as the foundation for the practical forecaster implementation, the ForecastICU algorithm, which we will present in the next section.

3.1. Mathematical model

Given a *probability space* $(\Omega, \mathcal{F}, \mathbb{P})$, and a measurable space (\mathcal{S}, Ψ) , a stochastic process $\mathbf{X}_{\mathcal{T}}$ comprises a collection of \mathcal{S} -valued random variables on Ω indexed by a totally ordered set \mathcal{T} . That is, $\mathbf{X}_{\mathcal{T}} = \{X_t\}_{t \in \mathcal{T}}$, where every X_t is an \mathcal{S} -valued random variable on Ω . Given the probability space $(\Omega, \mathcal{F}, \mathbb{P})$, a *filtration* $\{\mathcal{F}_t, t \in \mathcal{T}\}$ is a (weakly) increasing collection of σ -algebras on Ω . The filtration \mathcal{F}_t is always bounded above by \mathcal{F} , i.e. $\mathcal{F}_t \subseteq \mathcal{F}$. The *stopping time* τ_s of the process $\mathbf{X}_{\mathcal{T}}$ is a random variable $\tau_s : \Omega \rightarrow \mathcal{I}$, where \mathcal{I} is an ordered index set, e.g. $\mathcal{I} = [0, \infty)$. A stopping time τ_s satisfies that $\{\omega \in \Omega : \tau_s(\omega) \leq t\} \in \mathcal{F}_t, \forall t \in \mathcal{I}$. The stopping time of a process is decided by some arbitrary *stopping rule*. The *hitting time* τ_h of a process $\mathbf{X}_{\mathcal{T}}$ is the first time it hits a certain value, i.e. $\tau_h(\eta) = \inf\{t \in \mathbb{R} | X_t \geq \eta\}$.

We consider every patient’s physiological data stream as a stochastic process with respect to the space $(\Omega, \mathcal{F}, \mathbb{P})$. For instance, the blood pressure measurements stream can be viewed as a stochastic process $\mathbf{X}_{\mathcal{T}}$ that is observed starting from the time the patient was hospitalized. The stopping time τ_s of such a process is the time at which the physician takes a decision regarding the patient: the decision can be either discharging the patient or admitting her to the ICU; in both cases the patient is not monitored further. The stopping rule that determines such a stopping time is simply the physician’s criteria for ICU admission or discharging of hospitalized patients, which depends on how physicians interpret the physiological data stream.

Patients belong to two categories: stable patients who should be discharged, and clinically deteriorating patients who should be admitted to the ICU. We assume that the null hypothesis \mathcal{H}_0 is the hypothesis that the patient is stable, whereas the alternative hypothesis \mathcal{H}_1 is the hypothesis that the patient is clinically deteriorating. Depending on whether the true hypothesis is \mathcal{H}_0 or \mathcal{H}_1 , a physiological data stream $\mathbf{X}_{\mathcal{T}}$ will have a different joint distribution for its data samples with respect to the probability space $(\Omega, \mathcal{F}, \mathbb{P})$. We denote the families of finite-dimensional distributions of the physiological streams under the null and alternative hypotheses as \mathbb{P}_0 and \mathbb{P}_1 respectively, i.e. $\{X_{\tau}\}_{\tau=0}^t | \mathcal{H}_m \sim \mathbb{P}_m^t, m \in \{0, 1\}$, and \mathbb{P}_m is the family of distributions \mathbb{P}_m^t for all admissible values of t (e.g. maximum time a patient can stay hospitalized).

3.2. The Forecaster

3.2.1. FORMAL DEFINITION

We formally define the *forecaster* as a *belief system* that carries out the mapping $B_t : (\mathcal{F}_t, \mathbb{Q}_0, \mathbb{Q}_1) \rightarrow [0, 1]$, i.e. a map from a filtration to a *belief* about the monitored patient being clinically deteriorating, where \mathbb{Q}_0 and \mathbb{Q}_1 are

estimates of the families of finite-dimensional distributions \mathbb{P}_0 and \mathbb{P}_1 . We say that the belief system is *truthful* if $\mathbb{Q}_m = \mathbb{P}_m, m \in \{0, 1\}$, and we say that it is *non-truthful* otherwise, i.e. the estimated densities \mathbb{Q}_m of a non-truthful belief system have a non-zero distance from \mathbb{P}_m with respect to any probability metric. The function B_t (which we will also write as $B_t(\mathcal{H}_1 | \mathcal{F}_t)$) is denoted as the *belief function*, whereas the sequence $\{B_t\}_{t \in \mathcal{T}}$ is called the *belief process*. Informally, the forecaster is endowed with some information about the physiological stream generative process, encoded in the distributions \mathbb{Q}_0 and \mathbb{Q}_1 , and it accumulates information over time as it observes the monitored patient’s physiological streams, and builds a belief about her clinical status. Intuitively, the accuracy and timeliness of the forecaster depend on the quality of the estimates \mathbb{Q}_0 and \mathbb{Q}_1 , and the way the mapping B_t is implemented given the filtration \mathcal{F}_t .

3.2.2. STRUCTURE OF THE FORECASTER

The implementation of the forecaster’s belief system demands two basic modules: a *density estimation algorithm* \mathcal{A}^D , which finds “good” estimates \mathbb{Q}_0 and \mathbb{Q}_1 in an offline manner, and a *belief function* B_t , which updates the forecaster posterior belief in real-time in response to the information extracted from the monitored data streams. Formally, given that the forecaster has access to a dataset \mathbf{X}^{ref} comprising recorded data streams of N reference patients in the EHR who are labeled as being admitted to ICU or discharged, the density estimation algorithms is a mapping $\mathcal{A}^D : \mathbf{X}^{ref} \rightarrow (\mathbb{Q}_0, \mathbb{Q}_1)$, and the belief function is a real-time mapping $B_t : (\mathcal{F}_t, \mathcal{A}^D(\mathbf{X}^{ref})) \rightarrow [0, 1]$.

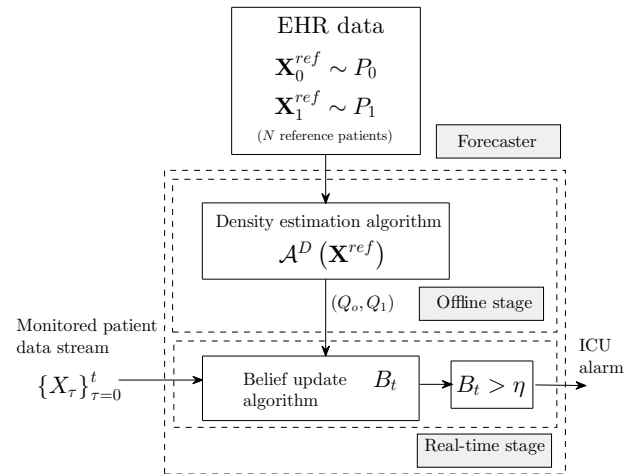


Figure 1. Schematic for a forecaster that learns to issue ICU admission alarms from the EHR data.

3.2.3. ALARM STRATEGY AND PERFORMANCE

The problem of (timely) predicting ICU admissions can be thought of as being equivalent to an optimal stopping problem, or a problem of *learning an unknown stopping rule* of a stochastic process. Thus, not only does the forecaster face uncertainty in the true hypothesis, but also it is uncertain about when will the process stop; the forecaster needs to figure out the true hypothesis before the process stops with a reasonable amount of time.

ForecastICU adopts a “threshold type” *alarm strategy*: it prompts an alarm for an ICU admission whenever its belief process first crosses a predefined threshold η . Thus, the optimal stopping time problem boils down to finding the optimal threshold that the belief process should hit before an ICU alarm is issued. Selection of the threshold η is aimed to maximize a *clinical value function* that comprises a set of accuracy and timeliness measures.

Three measures of performance are clinically relevant to the ICU prognostic setting: timeliness, sensitivity and precision (Mokart et al., 2013) (Prytherch et al., 2010) (Alam et al., 2014). Let τ_s be the stopping time of the patient’s physiological data stream, and let $\tau_h(\eta)$ be the hitting time of the belief process $\{B_t\}_{t \in \mathbb{N}}$ given a threshold η . The sensitivity of the forecaster, which is also known as the *true positive rate* (TPR), is given by

$$TPR = \frac{\mathbb{P}(\tau_h(\eta) < \tau_s | \mathcal{H}_1)}{\mathbb{P}(\tau_h(\eta) < \tau_s | \mathcal{H}_1) + \mathbb{P}(\tau_h(\eta) > \tau_s | \mathcal{H}_1)}, \quad (1)$$

whereas the precision, which is also known as the *positive predictive value* (PPV), is given by

$$PPV = \frac{\mathbb{P}(\tau_h(\eta) < \tau_s | \mathcal{H}_1)}{\mathbb{P}(\tau_h(\eta) < \tau_s | \mathcal{H}_1) + \mathbb{P}(\tau_h(\eta) < \tau_s | \mathcal{H}_0)}. \quad (2)$$

Finally, the timeliness of the forecaster $T_p(\eta)$ is the average time interval between the hitting time of the belief process and the stopping time of the physiological data stream for clinically deteriorating patients, which is formally given by

$$T_p(\eta) = \mathbb{E}[\tau_s - \tau_h(\eta) | \tau_h(\eta) < \tau_s, \mathcal{H}_1]. \quad (3)$$

The selection of the threshold value η should balance the trade-off between accuracy (in terms of TPR and PPV) and timeliness; intuitively, one expects that low threshold values would lead to more timely but less accurate decisions, and vice versa. Formally, we define a general clinical reward function $g(\{X_\tau\}_{\tau=0}^{\tau_h(\eta)})$ that quantifies the overall performance in terms of PPV, TPR and T_p as a function of the alarm strategy η . The ICU prognostic optimal stopping problem is equivalent to finding an optimal alarm strategy (a threshold η^*) that maximizes a *clinical value function* $V(\eta)$ as follows

$$V^* = \sup_{\eta \in [0, 1]} \mathbb{E}_{\mathbb{P}} \left[g(\{X_\tau\}_{\tau=0}^{\tau_h(\eta)}) \mathbf{1}_{\{\tau_h(\eta) < \tau_s\}} \right], \quad (4)$$

where the optimal clinical value function is $V^* = V(\eta^*)$. Problem (4) is challenging for that \mathbb{P} is unknown to the forecaster, and τ_s is random. The optimal solution to (4) would balance the value of information in the physiological stream $\{X_\tau\}_{\tau=0}^{\tau_h(\eta)}$ and the TPR, PPV and T_p which are reflected in the reward function g . Fig. 1 illustrates the structure of the forecaster: its offline and real-time components, and its alarm strategy.

3.3. Bayesian Learning

The forecaster’s alarm strategy is a threshold strategy on its belief process. The forecaster builds the belief process given observed information using a Bayesian learning approach: it updates its posterior belief in response to new observations of the physiological data streams as follows

$$B_t(\mathcal{H}_1 | \mathcal{F}_t) = \frac{\mathbb{Q}_1(\{X_\tau\}_{\tau=0}^t | \mathcal{H}_1) \mathbb{Q}_1(\mathcal{H}_1)}{\sum_{i \in \{0,1\}} \mathbb{Q}_i(\{X_\tau\}_{\tau=0}^t | \mathcal{H}_i) \mathbb{Q}_i(\mathcal{H}_i)},$$

where $\{\mathbb{Q}(\mathcal{H}_m)\}_{m \in \{0,1\}}$ are the forecaster’s estimates of the discrete priors. The forecaster cannot compute the optimal threshold in (4) since the belief system is generally non-truthful. Thus, the accuracy and timeliness of the forecaster (the value function $V(\eta)$) will be affected by the truthfulness of its belief system. In the following Theorem, we link the truthfulness of the forecaster’s belief system to the timeliness and accuracy of the alarm strategy in terms of the sample complexity. The proof of this Theorem is given in the Appendix.

Theorem 1 (Probably approximately correct and timely ICU alarm strategies) *For every $(\epsilon, \delta) \in [0, 1]^2$, there exists a polynomial function $N^*(\epsilon, \delta) = \text{poly}(\frac{1}{\delta}, \frac{1}{\epsilon})$ and a density estimation algorithm \mathcal{A}^D , such that for every dataset with $N > N^*(\epsilon, \delta)$ reference patients, we have that $\mathbb{P}(|V(\eta) - V^*| < \epsilon) \geq 1 - \delta$, where V^* is optimal clinical value function of a truthful belief system that has access to the optimal threshold η^* .*

Theorem 1 says that one can compute a sample complexity (the number of reference patients) that achieves an ϵ -optimal clinical value function (compared to an “oracle” belief system that has access to the true data streams’ distributions), with an arbitrary level of confidence $1 - \delta$. That is, the ICU prognosis problem using the proposed forecaster structure is *learnable* in the “probably approximately correct and timely” sense. This motivates the construction of a practical forecaster algorithm in the next section.

4. The ForecastICU Algorithm

Practical implementation of the forecaster is confronted with several challenges. First, not all the monitored physiological streams are relevant to the ICU prognostic decision.

Second, the data streams are not sampled at the same rate. Finally, a crucial aspect of the temporal physiological data is that the data streams of the reference patients are neither synchronized with each other nor synchronized with the incoming patients since each patient is hospitalized in a different clinical status.

As shown in Fig. 1, designing a forecaster entails designing a density estimation algorithm $\mathcal{A}^D(\mathbf{X}^{ref})$ and a belief updating procedure. In the following two subsections, we propose a design for both modules as the two building blocks of ForecastICU.

4.1. The offline density estimation algorithm $\mathcal{A}^D(\mathbf{X}^{ref})$

The algorithm $\mathcal{A}^D(\mathbf{X}^{ref})$ takes as an input a dataset \mathbf{X}^{ref} with N reference patient entries, associated with each patient a set of L recorded physiological data streams, and retrieves the estimated finite-dimensional distributions $(\mathbb{Q}_0, \mathbb{Q}_1)$. We denote the i^{th} reference patient by $\mathbf{X}_{(i)}^{ref}$, where $\mathbf{X}_{(i)}^{ref}$ is an $L \times K_i$ matrix, with K_i being the length of the longest data stream associated with reference patient i . We denote the overall hospitalization period of patient i as $T_{(i)}^H$. The algorithm implements the following four steps: **a- Non-causal alignment of reference patients’ data streams** : Since the reference patients are hospitalized with different states, ForecastICU needs to align their data streams prior to applying density estimation. The stopping time of the data streams usually follows a rational, consistent stopping rule decided by the physicians, and such a rule serves as a marker for the clinical status of the patient at the ICU admission time or discharge time. Therefore, ForecastICU aligns the patients’ data streams in a non-causal fashion: it views the stopping times of all streams in the dataset \mathbf{X}^{ref} as the reference time ($t = 0$) of the corresponding stochastic processes. We use the notation $\mathbf{X}_{(i)}^{ref}(m, n)$ to refer to the data sample of the m^{th} data stream at n time steps ahead of the stopping time ($t = 0$).

b- Interpolating under-sampled data streams : The different data streams for the same patient can be sampled with different sampling rates. Thus, the algorithm $\mathcal{A}^D(\mathbf{X}^{ref})$ would encounter missing values at some points in time when carrying out the joint density estimation across the features. To reconstruct the missing samples in under-sampled data streams, we use an *interpolation function* $h \left(\left\{ \mathbf{X}_{(i)}^{ref}(m, n) \right\}_{n=0}^{K-1}, T_s \right) : \mathbb{R}_{K \times 1} \times \mathbb{R} \rightarrow \mathbb{R}_{\lfloor \frac{KT_s}{T_{(i)}^H} \rfloor \times 1}$, i.e. a function that interpolates samples of a data stream and retrieves a stream with a sampling period of T_s instead of $\frac{T_{(i)}^H}{K}$. ForecastICU uses *cubic spline* interpolation function that we denote by h_{spline} . We denoted the interpolated dataset as $\tilde{\mathbf{X}}^{ref} = h_{spline}(\mathbf{X}^{ref})$, where all the data streams in $\tilde{\mathbf{X}}^{ref}$ are up-sampled with the sam-

pling rate of the most frequently sampled data stream.

c- Selecting relevant data streams : ForecastICU applies correlation feature selection (CFS) algorithm on the interpolated dataset $\tilde{\mathbf{X}}^{ref}$ (with minimum-redundancy-maximum relevance (mRMR) (Yu & Liu, 2003)) to discover the relevant temporal data streams for forecasting ICU patients. We denote the data stream selection operation as $\tilde{\mathbf{Y}}^{ref} = CFS(\tilde{\mathbf{X}}^{ref})$, where $\tilde{\mathbf{Y}}^{ref}$ is a dataset with N reference patients, associated with each a set of $R \leq L$ data streams, where R is the number of discovered relevant data streams.

d- Parametric density estimation : We estimate the finite-dimensional distributions $(\mathbb{Q}_0, \mathbb{Q}_1)$ based on the processed dataset $\tilde{\mathbf{Y}}^{ref}$ as follows. Since the dataset is labeled, we separate the reference patients into clinically deteriorating patients with processed data streams $\tilde{\mathbf{Y}}_1^{ref}$, and clinically stable patients with data streams $\tilde{\mathbf{Y}}_0^{ref}$. Let N_0 and N_1 be the number of entries in $\tilde{\mathbf{Y}}_0^{ref}$ and $\tilde{\mathbf{Y}}_1^{ref}$ respectively, and K_0 and K_1 be the average lengths of the data streams in $\tilde{\mathbf{Y}}_0^{ref}$ and $\tilde{\mathbf{Y}}_1^{ref}$ respectively. Let $(\mathbb{Q}_0^t, \mathbb{Q}_1^t)$ be the joint density functions of all the data samples for the R data streams selected by CFS between the (average) hospitalization time K_0 (or K_1) and the time instance that is t steps away from the stopping time. We approximate the finite-dimensional distributions $(\mathbb{Q}_0^t, \mathbb{Q}_1^t)$ as Multi-variate Gaussian distributions, and fit their mean and covariance parameters, i.e. $\mathbb{Q}_m^t \sim \mathcal{N}(\hat{\mu}_m^t, \hat{\Sigma}_m^t)$, $m \in \{0, 1\}$ where

$$[\hat{\mu}_m^t(j)]_{j=1}^R = \frac{1}{N_m} \sum_{i=1}^{N_m} \tilde{\mathbf{Y}}_{(i),m}^{ref}(j, t), \quad (5)$$

$$[\hat{\Sigma}_m^t]_{k,l} = \frac{1}{N_m - 1} \sum_{i=1}^{N_m} \tilde{\mathbf{Y}}_{(i),m}^{ref}(k, t) \tilde{\mathbf{Y}}_{(i),m}^{ref}(l, t), \quad (6)$$

where $k, l = 1, \dots, R$, and $\tilde{\mathbf{Y}}_{(i),m}^{ref}(k, t) = \tilde{\mathbf{Y}}_{(i),m}^{ref}(k, t) - \frac{1}{N_m} \sum_{w=1}^{N_m} \tilde{\mathbf{Y}}_{(w),m}^{ref}(k, t)$. Thus, the output of $\mathcal{A}^D(\mathbf{X}^{ref})$ is a family of estimated finite-dimensional distributions $(\mathbb{Q}_0, \mathbb{Q}_1)$, i.e. a set of Multi-variate Gaussian densities for different values of the elapsed hospitalization period which is able to capture temporal correlation. The density estimation algorithm passes the density estimates to the belief function which runs in real-time and applies Bayesian learning using $(\mathbb{Q}_0, \mathbb{Q}_1)$ as priors with which it updates its posterior beliefs.

4.2. The real-time belief updating algorithm

The belief function of ForecastICU is updated at time t given a the data stream $\{X_\tau\}_{\tau=t_0}^t$ (or generally, the information \mathcal{F}_t available at time t) of the incoming patient as follows $B_t(\mathcal{H}_1 | \mathcal{F}_t) = \mathbb{Q}_1(\mathcal{H}_1 | \{X_\tau\}_{\tau=t_0}^t)$. However, since ForecastICU faces uncertainty about the current clinical status (unknown the stopping time), the reference time

t_o of the stream $\{X_\tau\}_{\tau=t_0}^t$ with respect to the estimated densities $(\mathbb{Q}_0, \mathbb{Q}_1)$ is unknown. Therefore, ForecastICU estimates the belief function as

$$\frac{N_1 \mathbb{Q}_1(\{X_\tau\}_{\tau=t_0}^t | \mathcal{H}_1, T_1^*(t))}{N_0 \mathbb{Q}_0(\{X_\tau\}_{\tau=t_0}^t | \mathcal{H}_0, T_0^*(t)) + N_1 \mathbb{Q}_1(\{X_\tau\}_{\tau=t_0}^t | \mathcal{H}_1, T_1^*(t))}$$

where $t_o^m = t - T_m^*(t) + 1$, and $T_m^*(t)$ is the "most likely" estimate for the time remaining until the stopping time of the process, which is simply given by

$$T_m^*(t) = \arg \max_{\tau} \mathbb{Q}_m(\{X_k\}_{k=\tau-t+1}^{\tau} | \mathcal{H}_m). \quad (7)$$

The estimates are refined subsequently as additional measurements are observed.

In order to rule out drastic fluctuations and spikes in the belief function, which may result from a belief system constructed from a small data set, we apply a smoothing phase for the belief process via a simple moving average filter of length W . The smoothed belief function is given by $\tilde{B}_t = \frac{1}{W} \sum_{\tau=t-W}^t B_\tau$. ForecastICU prompts an ICU alarm whenever \tilde{B}_t exceeds a threshold η . In the following subsection, we show how the threshold η is set.

4.3. The alarm strategy

ForecastICU follows a threshold-type alarm strategy: it prompts an ICU alarm whenever the smoothed belief process \tilde{B}_t hits a threshold η . The fundamental trade-off that ForecastICU balances is the one between the timeliness of a prediction and its accuracy. We control such a trade-off via two parameters: the threshold η and the size of the smoothing filter W . In this case, the alarm strategy is defined by a richer set of parameters (η, W) rather than being solely determined by η . The threshold η controls to the confidence in the issued alarms, and W controls the stability of the belief function's fluctuations, and hence the system's sensitivity. In order to select desirable values for η and W , we define a clinical value function $V(\eta, W)$ as follows

$$V(\eta, W) = \alpha TPR(\eta, W) + (1 - \alpha) T_p(\eta, W) \quad (8)$$

where $\alpha \in [0, 1]$ is a parameter that balances the preferences over timeliness and sensitivity. By setting a constraint γ on the system's precision (which can be determined by physicians), ForecastICU computes the alarm strategy by solving the following optimization problem:

$$\begin{aligned} & \max_{\eta, W} && V(\eta, W) \\ & \text{subject to} && PPV(\eta, W) \geq \gamma. \end{aligned} \quad (9)$$

In the off-line stage, in which the complexity is less crucial, the computational complexity is $\mathcal{O}(N^2 D^3 W^3)$, where N is the number of patients, D is the number of features, and

W is the window size. In the real-time stage, the computational complexity is $\mathcal{O}(ND^2W^2)$. Therefore, the entire algorithm has cubic polynomial computation complexity regarding the number of patients.

5. Experiments

5.1. Data Description

ForecastICU was applied to a population of patients admitted to UCLA Ronald Reagan Medical Center. The population considered is fairly homogeneous: most patients were diagnosed with leukemia, lymphoma, multiple myelomas and other hematologic malignancies. The majority of the patients were receiving chemotherapy, allogeneic stem cell transplantation or autologous stem cell transplantation. We have chosen this particular population of patients for our experiments because these patients were receiving treatments (according to their diagnosis) that cause severe immunosuppression during their hospitalization, placing them at an extreme risk of developing a clinical deterioration, which requires ICU admission. This is extremely important as delays in ICU admission in these patients are associated with increased mortality and morbidity. (Liu et al., 2012) (Louay Mardini, 2012) (Young et al., 2003) (Mokart et al., 2013).

The patients' clinical features comprise 18 temporal physiological data streams which are described in the Appendix. Modeling these data streams as Gaussian model was validated by a Kolmogorov-Smirnov goodness-of-fit test (refer to the supplementary material). Each patient's data stream is associated with a binary label: either discharge (DIS) or ICU Admission (ICU). The total number of patients is 1065, and the number of patients admitted to ICU is 101 (9.48%). The hospitalization period ranges from 4 to 2062 hours (85 days and 22 hours).

5.2. Experiments Setup

We compare the performance of ForecastICU in terms of timeliness, PPV and TPR with respect to four state-of-the-art machine learning techniques: logistic regression (LR), support vector machines with radial based kernel (SVMs), regularized logistic regression with Lasso (LASSO) and random forest. The benchmarks operate by sequentially classifying the ICU patients as additional measurements are observed over time. Feature values are normalized by the maximum value of each physiological stream, and feature selection was applied to the benchmarks as well. The hyper-parameters of every benchmark was empirically optimized based on equation (9).

Performance measures were computed via 10-fold stratified cross-validation: we run 10 independent cross-

validations and report the average as the final performance. Clinicians responsible of ICU admissions in the UCLA Ronald Regan medical Center from which we extracted the data indicated that ICU alarms would be most helpful if they are at least 4 hours earlier than the standard time an ICU admission decision would be taken by the clinician, in order to provide sufficient time to safely enact the transfer and to potentially correct the cause of the underlying clinical deterioration.

5.3. Experiment Results

5.3.1. PERFORMANCE COMPARISONS

ForecastICU consistently outperforms the other benchmark algorithms with respect to the PPV for every given value of the TPR as shown in Fig. 2 and Table 1. For instance, for a TPR of 50%, ForecastICU achieves a PPV of 54.7%, which is 5.1% better than the best benchmark algorithm (random forest). Moreover, given a PPV of 40%, Forecast ICU achieves a TPR of 68.1%, which is 7.8% better than the best benchmark algorithm, with 8.15% gain in terms of the area under curve (AUC) as well. Note that while for different ranges of TPR and PPV, the best benchmark algorithm changes, ForecastICU is consistently outperforming that best benchmark.

Our algorithm offers a consistent improvement in terms of the rate of "false alarms" and "true alarms" as compared to the benchmarks; this is crucial in a practical ICU setting as studies suggest that the clinicians' response to an alarm is related to their confidence in the signal, which is directly related to how often the alarm correctly warns the clinician of impending danger (Cvach, 2012) (James P Bliss, 2010) (James P Bliss, 2007). For instance, at a TPR of 50%, ForecastICU is the only algorithm among those being compared that could correctly alarm for an ICU admission with an accuracy above 50% whereas other benchmarks perform worse than what a physician could do with a simple coin flip. The p -value of the hypothesis test that compares ForecastICU and the best algorithm is less than 0.01. Moreover, our algorithm can display the belief threshold for the computed alarm strategy, which itself is a natural measure for the "likelihood" of the alarm being truthful, and is thus a measure of confidence in the issued alarms.

5.3.2. TRADE-OFF BETWEEN TIMELINESS AND ACCURACY

Fig. 3 illustrate the trade-off between prediction time and its accuracy. While the performance of all algorithms naturally degrades as the ICU alarms are prompted at earlier times, we see that ForecastICU consistently outperforms all the other benchmark algorithms regarding both TPR and PPV for all alarm times (every level of timeliness). For in-

Table 1. Accuracy of ICU alarms for ForecastICU and the benchmark algorithms.

Algorithms	TPR(%)	PPV(%)
ForecastICU	50.2 ± 1.45%	54.7 ± 1.73%
Logistic Regression	50.8 ± 1.77%	39.7 ± 1.95%
Lasso Regularization	50.7 ± 2.01%	42.5 ± 1.98%
Random Forest	51.1 ± 2.03%	49.6 ± 1.55%
SVMs	50.5 ± 2.12%	29.8 ± 1.55%




Figure 2. Trade-off between TPR and PPV

stance, the TPR of ForecastICU is 59.2% with a PPV of 30.3% for a 12-hour early prediction with respect to the actual physician-determined ICU admission event. This represents a gain of 3.4% with respect to the best benchmark algorithm, which in this case is the LASSO regularization. Fig. 3 also shows that ForecastICU can consistently predict the ICU admission earlier than the best benchmark algorithm, raising the alarm around 2 to 3 hours earlier. It can also predict ICU admissions 9 hours before a physician’s decision for a PPV of 40% and a TPR of 50%.

6. Clinical significance of ForecastICU

The clinical significance of ForecastICU is not only limited to the performance gains, but it also extends to its ability to handle a versatile clinical value function, which can assist clinicians in managing the ICU admission procedure. For instance, given a TPR determined by the clinician, ForecastICU is able to warn the clinician earlier and provide a more confident signal than other existing machine learning algorithms, thus providing the busy clinician with a safety net for patient care by giving them sufficient time to intervene at an earlier time in order to prevent clinical deterioration. Moreover, given that the value of PPV is related




Figure 3. Trade-off between PPV and the prediction time (TPR is fixed at 50%).

(*: Timeliness of ForecastICU alarms compared to the best benchmark algorithm)

to the confidence that the clinician has the alarm, the clinician’s confidence with ForecastICU would be higher than other off-the-shelf benchmark algorithms regardless of the sensitivity (TPR) that the clinician decides to set. The PPV gains achieved by ForecastICU also imply a decrease in the number of patients that are falsely identified as needing ICU admission, which would reduce the reverse medical effects and unnecessary costs (Dasta Joseph F, 2005).

Table 2. Number of false alarm per one true alarm for different values of the TPR.

	TPR = 0.6	TPR = 0.55	TPR = 0.5
ForecastICU	1.14	0.96	0.84
Rothman index	3.11	2.67	2.34

Potential gains achievable by ForecastICU upon its actual deployment in a hospital ward is evaluated by constructing a comparison with the Rothman index, the currently deployed risk assessment technology in most wards (Rothman et al., 2013). We report a 22.3% gain in the AUC achieved by ForecastICU compared to that achieved by Rothman index. Moreover, in Table 2 we demonstrate the number of false alarms per one true alarm for both ForecastICU and the Rothman index at different settings of the TPR. At a TPR of 50%, ForecastICU leads to only 0.84 false alarms for every 1 true alarm, whereas the Rothman index lead to 2.34 false alarms per true alarm, i.e. the rate of the false alarms is higher than that of the true ones. Thus, ForecastICU can ensure more confidence in its issued ICU alarms, which would mitigate alarm fatigue and enhance a hospital’s resource utilization (Cvach, 2012)(James P Bliss, 2010).

Acknowledgements

The authors would like to thank the reviewers for their helpful comments. The research presented in this paper was supported by Air Force Dynamic Data-Driven Application Systems (DDDAS) Program.

References

- Alam, N, Hobbelink, EL, van Tienhoven, AJ, van de Ven, PM, Jansma, EP, and Nanayakkara, PWB. The impact of the use of the early warning score (ews) on patient outcomes: a systematic review. *Resuscitation*, 85(5):587–594, 2014.
- Bruijns, Stevan R, Guly, Henry R, Bouamra, Omar, Lecky, Fiona, and Lee, Wallis A. The value of traditional vital signs, shock index, and age-based markers in predicting trauma mortality. *Journal of Trauma and Acute Care Surgery*, 74(6):1432–1437, 2013.
- Buza, Krisztian, Nanopoulos, Alexandros, and Schmidt-Thieme, Lars. Time-series classification based on individualised error prediction. In *Computational Science and Engineering (CSE), 2010 IEEE 13th International Conference on*, pp. 48–54. IEEE, 2010.
- Churpek, Matthew M, Yuen, Trevor C, Park, Seo Young, Gibbons, Robert, and Edelson, Dana P. Using electronic health record data to develop and validate a prediction model for adverse outcomes on the wards. *Critical care medicine*, 42(4): 841, 2014.
- 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.
- Dasta Joseph F, McLaughlin Trent P, Mody Samir H Piech Catherine Tak. Daily cost of an intensive care unit day: the contribution of mechanical ventilation. *Critical care medicine*, 33(6): 1266–1271, 2005.
- 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.
- Ghassemi, Marzyeh, Pimentel, Marco AF, Naumann, Tristan, 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.
- Hüsken, Michael and Stagge, Peter. Recurrent neural networks for time series classification. *Neurocomputing*, 50:223–235, 2003.
- James P Bliss, Marica C Dunn. Behavioural implications of alarm mistrust as a function of task workload. *Ergonomics*, 43(9): 1283–1300, 2010.
- James P Bliss, Corey K Fallon, Nicolae Nica. The role of alarm signal duration as a cue for alarm validity. *Applied ergonomics*, 38(2):191–199, 2007.
- Liu, Vincent, Kipnis, Patricia, Rizk, Norman W, and Escobar, Gabriel J. Adverse outcomes associated with delayed intensive care unit transfers in an integrated healthcare system. *Journal of hospital medicine*, 7(3):224–230, 2012.
- Louay Mardini, Jed Lipes, Dev Jayaraman. Adverse outcomes associated with delayed intensive care consultation in medical and surgical inpatients. *Journal of critical care*, 27(6):688–693, 2012.
- Mokart, Djamel, Lambert, Jérôme, Schnell, David, Fouché, Louis, Rabbat, Antoine, Kouatchet, Achille, Lemiale, Virginie, Vincent, François, Lengliné, Etienne, Bruneel, Fabrice, 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.
- Neti, Chalapathy, Ebadollahi, Shahram, Hu, Jianying, Gotz, David, Sun, Jimeng, and Sow, Daby. Predicting patient's trajectory of physiological data using temporal trends in similar patients: A system for near-term prognostics. 2010.
- Pimentel, Marco AF, Clifton, David A, Clifton, Lei, Watkinson, Peter J, and Tarassenko, Lionel. Modelling physiological deterioration in post-operative patient vital-sign data. *Medical & biological engineering & computing*, 51(8):869–877, 2013.
- Prytherch, David R, Smith, Gary B, Schmidt, Paul E, and Featherstone, Peter I. Views towards a national early warning score for detecting adult inpatient deterioration. *Resuscitation*, 81(8):932–937, 2010.
- 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.
- Seker, Huseyin, Odetayo, Michael O, Petrovic, Dobrila, and Naguib, Raouf NG. A fuzzy logic based-method for prognostic decision making in breast and prostate cancers. *Information Technology in Biomedicine, IEEE Transactions on*, 7(2):114–122, 2003.
- Steyerberg, Ewout W, Eijkemans, Marinus JC, Harrell, Frank E, and Habbema, J Dik F. Prognostic modelling with logistic regression analysis: a comparison of selection and estimation methods in small data sets. *Statistics in medicine*, 19(8):1059–1079, 2000.
- Steyerberg, Ewout W, Eijkemans, Marinus JC, Harrell, Frank E, and Habbema, J Dik F. Prognostic modeling with logistic regression analysis in search of a sensible strategy in small data sets. *Medical Decision Making*, 21(1):45–56, 2001.
- Sun, Jimeng, Sow, Daby, Hu, Jianying, and Ebadollahi, Shahram. A system for mining temporal physiological data streams for advanced prognostic decision support. In *Data Mining (ICDM), 2010 IEEE 10th International Conference on*, pp. 1061–1066. IEEE, 2010.
- Tsien, Christine L and Fackler, James C. Poor prognosis for existing monitors in the intensive care unit. *Critical care medicine*, 25(4):614–619, 1997.

- Van Heeswijk, Mark, Miche, Yoan, Lindh-Knuutila, Tiina, Hilbers, Peter AJ, Honkela, Timo, Oja, Erkki, and Lendasse, Amaury. Adaptive ensemble models of extreme learning machines for time series prediction. In *Artificial Neural Networks-ICANN 2009*, pp. 305–314. Springer, 2009.
- Wei, Li and Keogh, Eamonn. Semi-supervised time series classification. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 748–753. ACM, 2006.
- Xing, Zhengzheng, Pei, Jian, and Philip, S Yu. Early classification on time series. *Knowledge and information systems*, 31(1): 105–127, 2012.
- Young, Michael P, Gooder, Valerie J, Bride, Karen, James, Brent, and Fisher, Elliott S. Inpatient transfers to the intensive care unit. *Journal of general internal medicine*, 18(2):77–83, 2003.
- 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.
- Zhang, Yang, Fong, Simon, Fiaidhi, Jinan, and Mohammed, Sabah. Real-time clinical decision support system with data stream mining. *BioMed Research International*, 2012, 2012.