DYNACARE: Dynamic Cardiac Arrest Risk Estimation

Joyce C. Ho[1], **Yubin Park**[1], **Carlos M. Carvalho**[2], **Joydeep Ghosh**[1]
[1] Department of Electrical and Computer Engineering
[2] McCombs School of Business
The University of Texas at Austin

Abstract

Cardiac arrest is a deadly condition caused by a sudden failure of the heart with an in-hospital mortality rate of $\sim 80\%$. Therefore, the ability to accurately estimate patients at high risk of cardiac arrest is crucial for improving the survival rate. Existing research generally fails to utilize a patient's temporal dynamics. In this paper, we present two dynamic cardiac risk estimation models, focusing on different temporal signatures in a patient's risk trajectory. These models can track a patient's risk trajectory in real time, allow interpretability and predictability of a cardiac arrest event, provide an intuitive visualization to medical professionals, offer a personalized dynamic hazard function, and estimate the risk for a new patient.

1 Introduction

Cardiac arrest is an abrupt cessation of heart function that prevents blood circulation. Disturbances in the electrical system of the heart may lead to abnormal heart rhythms, halting the pumping action of the heart. Common causes of cardiac arrest are ventricular tachycardia (irregular heartbeat caused by a fast heart rate), ventricular fibrillation (uncontrolled twitching of the heart muscles), asystole (sudden pause of heart muscle contractions), or pulseless electrical activity (no detectable heartbeat). For every 1000 hospital admissions, approximately 5 patients experience a cardiac arrest event with a mortality rate of $\sim 80\%$ (Sandroni et al., 2007). Studies have shown that $\sim 62\%$ of cardiac arrests could been prevented based on clinical evidence of deterioration 8 hours prior to the event (Hodgetts et al., 2002; Sandroni et al., 2007; Churpek et al., 2012). In addition, a quick response to cardiac arrest can decrease the mortality rate to 60% (Andréasson et al., 1998; Sandroni et al., 2004). However, the inability to correctly identify patients with sufficient intervention time limits the effectiveness of emergency response teams (Churpek et al., 2012). Therefore, accurate identification of at-risk patients is critical to minimizing the number of cardiac arrests and improving the survival rate.

The advent of electronic health records (EHR) has increased the availability of medical data. The Multiparameter Intelligent Monitoring in Intensive Care II (MIMIC-II) database is the most extensive and publicly available intensive care unit (ICU) resource. It was developed to support research in clinical decision support and critical care medicine (Saeed et al., 2011). Data was collected over 30,000 ICU patients during 2001 to 2007 from Boston's Beth Israel Deaconess Medical Center. The MIMIC-II database allows us to explore and evaluate models to estimate the risk of cardiac arrest over a large population of patients.

Recent research has focused on establishing early warning scores or criteria for predicting patients at high risk of experiencing a cardiac arrest. Many published physiologically-based criteria exist to detect patient deterioration and could be used to predict adverse outcomes (Smith and Wood, 1998). One set of early detection criteria used doctor or nurse concerns, respiratory rate, blood pressure, and temperature measurements to alert an emergency response team (Hodgetts et al., 2002). However the various criteria, including the Modified Early Warning Score (McBride et al., 2005) are based primarily on expert opinion and have limited scientific validation (Churpek et al., 2012). To address these shortcomings, Churpek et al. (2012) proposed the use of a scoring system derived from vital signs in the ward to detect clinical deterioration. Although scoring systems or activation criteria can identify high-risk patients, they are unable

Appearing in Proceedings of the 16^{th} International Conference on Artificial Intelligence and Statistics (AISTATS) 2013, Scottsdale, AZ, USA. Volume 31 of JMLR: W&CP 31.

to predict the time of cardiac arrest. These systems fail to capture temporal patterns in the physiological measurements. Kennedy and Turley (2011) suggested adding clinically relevant latent variables, trend features, and seasonality features to supplement the raw time series. Other approaches have involved searching for temporal patterns within the data (Batal et al., 2012; Wang et al., 2012). We propose an approach based on dynamic time series models common in economic forecasting to predict the time of cardiac arrest for high-risk patients.

This paper presents two dynamic cardiac arrest risk estimation (DYNACARE) models, variations of the dynamic stochastic volatility factor model proposed by Carvalho et al. (2011). We investigate the application of a semi-supervised framework to a dynamic non-linear regression model. The DYNACARE models (i) continuously track a patient's cardiac risk trajectory, (ii) allow interpretability and predictability of a cardiac arrest event, (iii) provide an intuitive visualization of a patient's cardiac arrest, (iv) deliver real-time results through a distributed implementation, (v) provide a dynamic hazard function unobtainable via traditional analysis, and (vi) generalize for any new patient.

Notation Preliminaries. Lowercase letters represent scalars, for example λ, r. Lowercase boldface letters, such as $\mathbf{y}, \boldsymbol{\mu}$, are vectors. Uppercase boldface letters correspond to matrices, for example $\boldsymbol{\Sigma}$. The subscript notation r_t represents the value of r at time t. $r_{1:t}$ is then the set of values from time 1 to time t.

2 DYNACARE Models

We model a patient's cardiac arrest trajectory (CAT) as a single latent factor, illustrated in Figure 1. The sequence of physiological measurements are a function of the patient's CAT. The simplest model, a general dynamic linear model with Kalman filter-forward steps and backward smoothing, was unable to fully capture the data. Figure 2 motivates the use of a stochastic volatility (SV) model as the variance of the risk residuals seemed to be auto-correlated.

The standard SV model assumes that the variance of returns on assets follows a latent stochastic process (Kim et al., 1998). Gibbs sampling can be used to explore the conditional posterior distribution of all the states (Kim et al., 1998). Additionally, it has been shown that a particle implementation of forward filtering - backward smoothing can be used to simulate the SV model (Doucet and Johansen, 2008). Carvalho et al. (2011) proposed a stochastic volatility factor model where the factors are driven by univariate SV

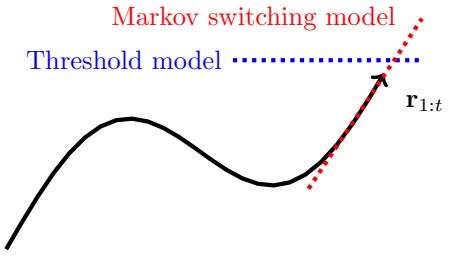

Figure 1: A sample cardiac risk trajectory and the temporal signature associated with the two DYNACARE models.

models.

Our general DYNACARE model extends the general dynamic linear model and assumes the variance of the risk trajectory is driven by a SV model. Sections 2.1 and 2.2 present two instantiations of the DYNACARE model. Equation block 1 illustrates our generalized model. In DYNACARE, r is the latent factor CAT, λ is the stochastic volatility term, and $\mathbf{y}$ is the set of observations with f unique measurement types.

$$\begin{aligned} \lambda_t &= \lambda_{t-1} + \delta_t & \delta_t &\sim N(0, k^2) \\ r_t &= \alpha_t + r_{t-1} + \varepsilon_t & \varepsilon_t &\sim N(0, \exp(\lambda_t)) \\ \mathbf{y}_t &= \boldsymbol{\mu} + \boldsymbol{\beta} r_t + \boldsymbol{\eta}_t & \boldsymbol{\eta}_t &\sim N(0, \boldsymbol{\Sigma}) \\ \boldsymbol{\Sigma} &= \mathrm{diag}(\sigma_1^2, \sigma_2^2 \cdots, \sigma_f^2) \end{aligned} \tag{1}$$

Standard particle smoothing approaches to the stochastic volatility model are insufficient for our model as the latent factor is unrelated to the cardiac arrest event. DYNACARE employs a semi-supervised framework to link the obtained latent factors to the rare event. Although we cannot ascertain the period in which a patient is healthy or if any unrecorded or unobserved cardiac arrest events transpired, our models incorporate the fact that we know a cardiac arrest event occurred at a specific time point. This information is utilized in the "backward smoothing" step of our particle filtering algorithm. Thus, DYNACARE provides interpretability of the latent factor as well as predictability of the cardiac arrest event.

The general DYNACARE algorithm combines the expectation maximization (EM) algorithm and particle smoothing. To prevent degeneracy, where a single unique particle approximates $p(r_{1:n}|\mathbf{y}_{1:T})$ for $n << T$, our algorithms use fixed-lag approximation. This leverages the forgetting properties of hidden Markov models such that for Δ large enough, $p(r_{1:n}|\mathbf{y}_{1:T}) \approx p(r_{1:n}|\mathbf{y}_{1:\min(n+\Delta,T)})$. For each patient, we use a model-specific particle smoother with fixed-lag approximation to estimate the latent variables. Model parameters are then obtained from the estimated latent variables. The process iterates un-

til convergence of the parameters occurs. The general framework is outlined in Algorithm 1.

Algorithm 1 General DYNACARE patient algorithm

for i=1 : M **do**
  Estimate $\hat{r}_{1:T}$ with model-specific particle smoother
  Learn $\boldsymbol{\beta}_i, \boldsymbol{\Sigma}_i$ given $\hat{r}_{1:T}$
end for
$\boldsymbol{\beta}, \boldsymbol{\Sigma}$ as average of last 10 samples

We present two specializations of the DYNACARE model that estimate risk of cardiac arrest for patients through a temporal CAT signature. The Markov switching model uses a gradient-based approach to define cardiac arrest. The second model, the threshold model, is a locality-based approach. Figure 1 shows the temporal signature associated with the different models.

2.1 Markov Switching Model

Markov switching model, also known as Markov switching multifractal, is a widely adopted modeling framework in financial econometrics to incorporate heterogeneous stochastic volatility (Calvet and Fisher, 2004). In DYNACARE Markov switching model (MSM), we assume two heterogeneous dynamics of the risk factor, namely a healthy and risky state. These states govern the gradient (or difference) of the observations. For MSM, cardiac arrest occurs when a patient is at the risky state. Exploratory data analysis confirmed abrupt gradient changes for some cardiac arrests and the risky state attempts to capture such movements. The transition probability from the healthy state (s_h) to the cardiac-arrest risky state (s_c) is given as p_{hc}, and from the risky state to the healthy state as p_{ch}. The stationary distribution of the two states is then $(\pi_h, \pi_c) = (\frac{p_{ch}}{p_{ch}+p_{hc}}, \frac{p_{hc}}{p_{ch}+p_{hc}})$.

Maximum likelihood (ML) estimation of these parameters would result in $p_{ch} \gg p_{hc}$, as the probability of cardiac arrest events is extremely rare amongst all the patients. However, this ML parametrization would drastically decrease the sensitivity of the model. We treat these transition parameters as knobs of the model that control the model sensitivity. Note that the probability of staying in the healthy state for time T before jumping to the risky state is $(1-p_{hc})^T p_{ch}$. If the desired cardiac arrest event notification time is within time period T_{min}, the parametrization should satisfy the condition $(1-p_{hc})^{T_{min}} p_{ch} > p_{\text{threshold}}$. For this work, we assume an equal stationary density $(\pi_h, \pi_c) = (0.5, 0.5)$ and $p_{hc} = 0.2$, but the settings should be changed depending on the objectives. Our experimental results, which are not provided due to page length constraints, show that higher values of p_{hc} results in both higher false positives and true positives. Moreover, low values of p_{ch} and p_{hc} impart inertia on the states.

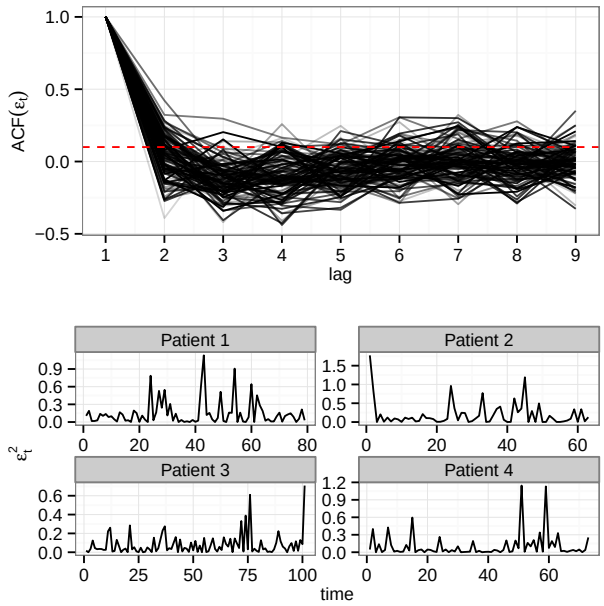

Figure 2: Auto-correlation of ε_t without the SV model (top), and ε_t^2 vs Time for randomly selected four patients (bottom). High-variance noise co-occur in a short time window.

The stochastic volatility (SV) model introduces another independent underlying stochastic process in addition to the two-state Markov process. Figure 2 shows the auto-correlation of ε_t^2 without the SV model. The result supports the use of time-varying variance. We assume that the stochastic volatility model in DYNACARE is wide-sense stationary, thus $\mathrm{E}[\lambda_t] = 1$, $\mathrm{Var}(\lambda_t) = k^2$, where the variance k^2 is sampled from a non-informative Inverse-Gamma prior distribution.

The risk trajectory in MSM is a function of these two underlying processes, u_t and λ_t. The stochastic volatility term λ_t not only models the auto-correlation among risk factor residuals (inter-correlation), but also captures individual differences in the risk residuals (intra-correlation). In other words, the variability of the risk factor varies from person to person, as well as from time to time. MSM can be formally written as follows:

$$
\begin{aligned}
&\lambda_t = \lambda_{t-1} + \delta_t && \delta_t \sim N(0, k^2) \\
&u_t \sim \text{MarkovChain}(u \mid u_{t-1}) && u_t \in \{s_h, s_c\} \\
&r_t = \alpha_{u_t} + \varepsilon_t && \varepsilon_t \sim N(0, \exp(\lambda_t)) \\
&\alpha_{u_t} \in \{\alpha_{s_h}, \alpha_{s_c}\},\ \alpha_{s_h} \neq \alpha_{s_c} && \\
&\Delta \mathbf{y}_t = \mathbf{y}_t - \mathbf{y}_{t-1} = \boldsymbol{\beta} r_t + \boldsymbol{\eta}_t && \boldsymbol{\eta}_t \sim N(0, \boldsymbol{\Sigma}) \\
&\boldsymbol{\Sigma} = \text{diag}(\sigma_1^2, \cdots, \sigma_f^2) &&
\end{aligned}
\tag{2}
$$

Table 1: Important MSM distributions

Joint distribution	$p(\mathbf{y}_{1:T}, r_{1:T}, u_{1:T}, \lambda_{1:T}, \boldsymbol{\beta}, \boldsymbol{\Sigma}) = p(\mathbf{y}_{1:T}\|r_{1:T}, \boldsymbol{\beta}, \boldsymbol{\Sigma})p(r_{1:T}\|u_{1:T}, \lambda_{1:T})p(u_{1:T}, \lambda_{1:T})p(\boldsymbol{\beta}, \boldsymbol{\Sigma})$
Filter forward	$p(r_{t+1}, u_{t+1}, \lambda_{t+1}\|r_t, u_t, \lambda_t, \mathbf{y}_{1:(t+1)}, \boldsymbol{\beta}, \boldsymbol{\Sigma})$ $\propto p(\mathbf{y}_{t+1}\|r_{t+1}, \boldsymbol{\beta}, \boldsymbol{\Sigma})p(r_{t+1}\|u_{t+1}, \lambda_{t+1})p(u_{t+1}\|u_t)p(\lambda_{t+1}\|\lambda_t)$
Backward smooth	$p(r_{1:T}, u_{1:T}, \lambda_{1:T}\|\mathbf{y}_{1:T}) = p(r_T, u_T, \lambda_T\|\mathbf{y}_T) \prod_{t=T-1}^{1} p(r_t, u_t, \lambda_t\|r_{(t+1):T}, u_{(t+1):T}, \lambda_{(t+1):T}, \mathbf{y}_{1:t})$

Figure 3 shows its graphical representation of the model.

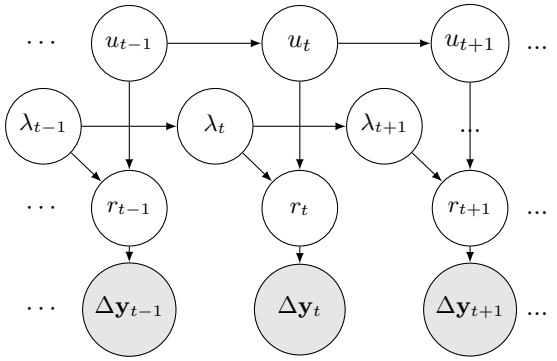

Figure 3: The graphical representation of the DYNACARE Markov switching model (MSM).

The joint distribution of MSM is described in Table 1. DYNACARE uses the EM algorithm to estimate the parameters $\boldsymbol{\beta}, \boldsymbol{\Sigma}$.

$$\{\hat{r}_{1:T}, \hat{u}_{1:T}, \hat{\lambda}_{1:T}\} \sim E[r_{1:T}, u_{1:T}, \lambda_{1:T}|\mathbf{y}_{1:T}, \boldsymbol{\beta}, \boldsymbol{\Sigma}] \quad (3)$$
$$\{\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\Sigma}}\} \sim \max p(\mathbf{y}_{1:T}, r_{1:T}, u_{1:T}, \lambda_{1:T}|\boldsymbol{\beta}, \boldsymbol{\Sigma})$$

Equation 3 can be efficiently simulated using a particle smoother, with steps detailed in Table 1. Partial knowledge that the cardiac arrest event occurred at the last time period is incorporated via the following:

$$p(u_T = s_c|\mathbf{y}_T) \approx 1 \quad (4)$$

Algorithm 2 illustrates the overall procedure of performing the MSM particle smoother.

2.2 Threshold Model

A threshold model is commonly used in toxicology to model the concept that doses above a certain level are dangerous, while anything below that is safe (Cox, 1987). Cardiac arrest is defined using a similar notion, where the event occurs when the CAT exceeds a specific value. In the DYNACARE threshold model (THR), the risk trajectory is a function of only one underlying process λ_t. The THR model can then be

Algorithm 2 MSM particle smoother

Draw $k^{(i)} \sim \Gamma^{-1}(\alpha_k, \beta_k)$
Draw $\lambda_0^{(i)} \sim \mathcal{N}(0, k^{(i)}), u_0^{(i)} \in \{0, 1\}, r_0^{(i)} \sim \mathcal{N}(0, 1)$
for $t = 1 : t_{\text{CA}}$ **do**
  for $\tau = t : \min(t + L, t_{\text{CA}})$ **do**
    Draw $\lambda_\tau^{(i)} \sim \mathcal{N}(\lambda_{\tau-1}^{(i)}, k^{(i)})$
    Draw $u_\tau^{(i)} \sim \text{MarkovChain}(u \mid u_{\tau-1}^{(i)})$
    Draw $r_\tau^{(i)} \sim \mathcal{N}(\alpha_{u_\tau^{(i)}}, \exp(\lambda_\tau^{(i)}))$
    $w_\tau^{(i)} \propto \exp(\frac{1}{2}(\mathbf{y}_\tau - \boldsymbol{\beta} r_\tau^{(i)})^\top \boldsymbol{\Sigma}^{-1}(\mathbf{y}_\tau - \boldsymbol{\beta} r_\tau^{(i)}))$
  end for
  for $\tau = \min(t + L, t_{\text{CA}}) : t$ **do**
    $w_{\tau-1}^{(i)} \propto w_\tau^{(i)} p(r_\tau^{(i)}|r_{\tau-1}^{(i)}, u_\tau^{(i)}, \lambda_\tau^{(i)}) p(u_\tau^{(i)}|u_{\tau-1}^{(i)}) p(\lambda_\tau^{(i)}|\lambda_{\tau-1}^{(i)})$
  end for
  $\hat{u}_t = \sum w_t^{(i)} u_t^{(i)}$
  $\hat{r}_t = \sum w_t^{(i)} r_t^{(i)}$
end for

written as:

$$\begin{aligned} \lambda_t &= \lambda_{t-1} + \delta_t & \delta_t &\sim N(0, k^2) \\ r_t &= r_{t-1} + \varepsilon_t & \varepsilon_t &\sim N(0, \exp(\lambda_t)) \\ \mathbf{y}_t &= \boldsymbol{\beta} r_t + \boldsymbol{\eta}_t & \boldsymbol{\eta}_t &\sim N(0, \boldsymbol{\Sigma}) \\ \boldsymbol{\Sigma} &= \text{diag}(\sigma_1^2, \sigma_2^2 \cdots, \sigma_f^2) \end{aligned}$$

Figure 4 shows the graphical representation of the proposed model.

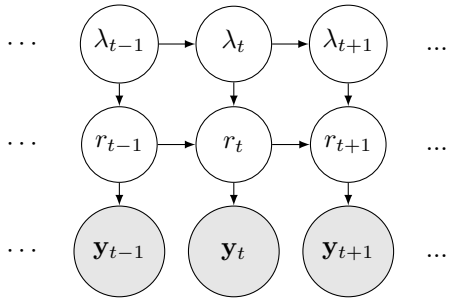

Figure 4: The graphical representation of DYNACARE threshold Model (THR).

For THR, cardiac arrest is defined as the point where the risk trajectory r_t exceeds a certain value θ. Furthermore, the model assumes that the risk trajectory is monotonically increasing in a time period (L) before cardiac arrest. We impose the following restrictions on r_t during the semi-supervised backward smoothing

stage of our algorithm:

$$\begin{aligned} r_t &\geq \theta & \forall t \geq & t_{\text{CA}} \\ r_t &\geq r_{t-1} & \forall t \geq & t_{\text{CA}} - L \end{aligned} \tag{5}$$

The EM algorithm to estimate the THR parameters $\boldsymbol{\beta}, \boldsymbol{\Sigma}$ is the same as in the MSM model (Equation 3). Additionally, the particle smoother for THR has an analogous form to the MSM particle smoother. The main difference between the two models lies in the incorporation of the partial knowledge about the cardiac arrest event. THR enforces Equation 5 through: (i) a penalty on the particle weights, and (ii) sampling $r_t^{(i)}$ from a truncated normal distribution to bound $r_t^{(i)}$ such that $r_{t-1}^{(i)} \leq r_t^{(i)}$ when $t \geq t_{\text{CA}} - L$. The penalty factor, $\rho \ll 1$, decreases the weight of particles that violate Equation 5.

$$\tilde{w}_{t_{\text{CA}}}^{(i)} = \begin{cases} w_{t_{\text{CA}}}^{(i)} & r^{(i)} \geq \theta \\ \rho w_{t_{\text{CA}}}^{(i)} & \text{otherwise} \end{cases}$$

The procedure for THR particle smoothing is detailed in Algorithm 3.

Algorithm 3 THR particle smoother

```
Draw k^(i) ~ Γ^{-1}(α_k, β_k)
Initialize λ_0^(i) ~ N(0, k^(i)), r_0^(i) ~ N(0, 1)
for t = 1 : t_CA do
  for τ = t : min(t + L, t_CA) do
    Draw λ_τ^(i) ~ N(λ_{τ-1}^(i), k^(i))
    Draw r_τ^(i) ~ { N(r_τ^(i), exp(λ_τ^(i))),                 τ < T − L
                   { TN(r_τ^(i), exp(λ_τ^(i)), r_τ^(i), ∞),    τ ≥ T − L
    w_τ^(i) ∝ exp((1/2)(y_τ − β r_τ^(i))^⊤ Σ^{-1} (y_τ − β r_τ^(i)))
    if r_τ^(i) < θ and τ = t_CA then
      w̃_τ^(i) = ρ w_τ^(i)
    else
      w̃_τ^(i) = w_τ^(i)
    end if
    w_τ^(i) = w̃_τ^(i)
  end for
  for τ = min(t + L, t_CA) : t do
    w_{τ-1}^(i) ∝ w_τ^(i) p(r_τ^(i) | r_{τ-1}^(i), λ_τ^(i)) p(λ_τ^(i) | λ_{τ-1}^(i))
  end for
  r̂_t = Σ w_t^(i) r_t^(i)
end for
```

3 DYNACARE Benefits

DYNACARE learns an individual patient's model parameters and estimates the cardiac arrest trajectory. In addition, it can model a new patient, deliver instantaneous results for a large patient population via distributed computing, and provide a personalized dynamic hazard function.

3.1 Algorithm Parallelization

The DYNACARE algorithm estimates the model parameters and cardiac arrest trajectory for each patient.

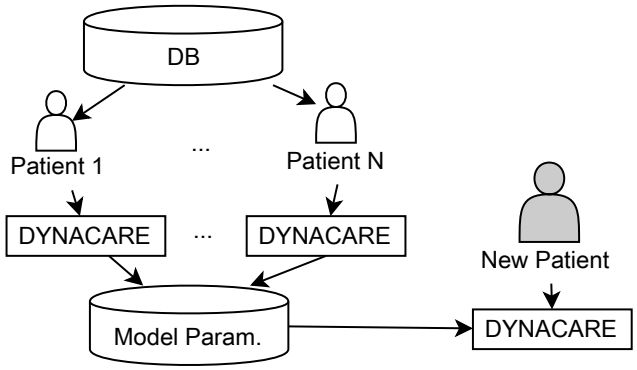

Figure 5: Diagram of the implemented distributed DYNACARE system. DYNACARE is embarrassingly parallelizable.

Consequently, the computation is distributed across multiple machines; each system tasked with learning an individual's parameters and CAT. A database is then used to store all the learned patient parameters to model new patients with insufficient number of observations. Furthermore, the particle smoother itself can be parallelized using a MapReduce framework. The "map" function takes the current risk, propagates it forward and calculates the weight based on the distribution. The "reduce" function renormalizes and resamples the new particle weights. Figure 5 illustrates the distributed systems diagram for DYNACARE.

3.2 Survival Analysis

Survival analysis defines a hazard function $(h(t))$, the instantaneous rate of failure at time t conditioned on survival up to t, as $h(t)dt = p(t < t_{\text{event}} < t + dt | t_{\text{event}} \geq t)$. The Cox proportional hazard model (Cox, 1972) is a popular semi-nonparametric model. Classical hazard models can not be used for this problem as the cardiac arrest event time cannot be aligned across patients, violating a major assumption of survival analysis. However, we will show that DYNACARE dynamically tracks the current state of the patient and provides a personalized dynamic hazard function $h(t)$.

The DYNACARE models assume the cardiac arrest trajectory is a wide-sense stationary random process. Without loss of generality, $E[r_t] = 0$ and $E[\lambda_t] = 0$. For a new patient without any observations, MSM assigns the probability of experiencing a cardiac arrest at time t as π_c. After observing a sequence of measurements $\Delta\mathbf{y}_{1:t}$, the model estimates $\hat{u}_t | \Delta\mathbf{y}_{1:t}$. Thus the probability of a cardiac arrest event at time t is the transition probability from $\hat{u}_t$ to s_c, providing a personalized hazard function that varies over time and

patient history.

$$h_{MSM}(t) = p(\hat{u}_t = s_c|\Delta \mathbf{y}_{1:t})$$

THR assigns the probability of a cardiac arrest event in the next time period for a new patient as $h_{THR}(t) = p(\varepsilon_t > \theta)$, $\varepsilon_t \sim N(0,1)$. As the model observes the patient's information $\mathbf{y}_{1:t}$, the new hazard function takes a dynamic form.

$$h_{THR}(t) = p(\hat{r}_t + \varepsilon_t > \theta \mid \mathbf{y}_{1:t}) = p(\varepsilon_t > \theta - \hat{r}_t \mid \mathbf{y}_{1:t})$$
$$= \int_{\theta-\hat{r}_t}^{\infty} \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{x^2}{2}\right) dx = 1 - \mathbf{\Phi}(\theta - \hat{r}_t)$$

$\mathbf{\Phi}$ represents the cumulative normal distribution in the equation above.

4 Experiment

4.1 Data

The study was conducted on adults (18+ years of age at time of admission) from the MIMIC-II database who had an asystole event. We focused on five measurement types: heart rate, respiratory rate, body temperature, diastolic blood pressure and systolic blood pressure. Data prior to cardiac arrest time was discretized into 4-hour bins starting when a patient has at least one observation per measurement. Additionally, we required each patient to have at least 40 discrete time slices ($\sim$ 6.5 days) to ensure sufficient data points.

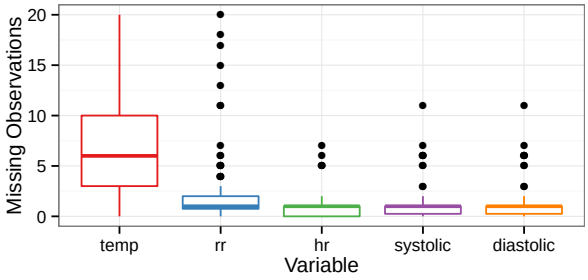

Figure 6: Boxplot summarizing the number of no measurement time slices per variable.

From 27,542 adult hospital admissions, there were 421 cardiac arrest patients with asystole. However only 108 of these patients met the minimum data requirements. On average, patients had 76 time slices with a standard deviation of 22. We assumed unobserved measurements denote the patient's status quo and employed the zero-order hold (Fialho et al., 2010), maintaining the last observed value. Figure 6 displays the number of missing observations for each measurement type per patient. Heart rate, diastolic blood pressure and systolic blood pressure were generally observed at every time slice. On the other hand, temperature measurements were not regularly measured every 4 hours. Figure 7 shows a plot of the last 100 time periods prior to cardiac arrest for a patient.

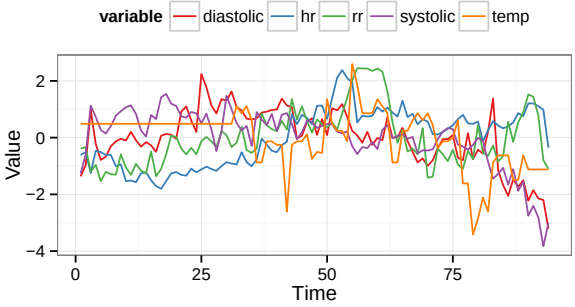

Figure 7: An example of a patient's normalized physiological measurements prior to cardiac arrest.

4.2 Evaluation Measure

Learned model parameters are utilized to estimate the risk of a new patient. An exploratory analysis of the learned parameter distribution showed an underlying hierarchical structure, which is illustrated in Figure 8. Model parameters are drawn from stratified learned-parameter samples based on a patient's age and gender, which we refer to as stratified bootstrapping (SB). Stratified bootstrapping is used as a computationally efficient alternative to modeling the hierarchical structure directly. Table 2 shows the number of patients per subgroup, or stratum.

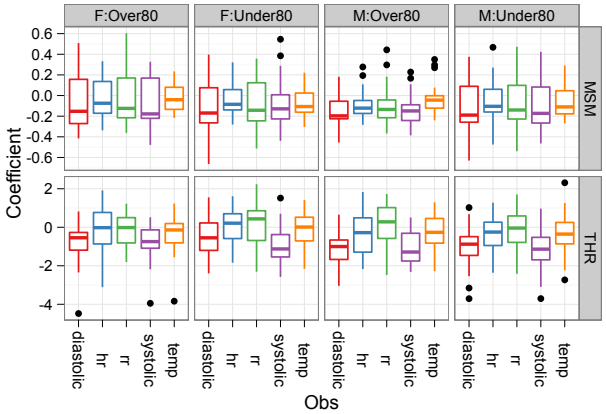

Figure 8: The distribution of beta parameters based on a patient's age group and gender. Each stratum exhibits a different mean and variance of the estimated parameters.

The predictive performance of the DYNACARE models, an unsupervised simple dynamic linear model, and

Table 2: Number of patients per strata

	Age under 80	Age over 80
Female	23	19
Male	49	17

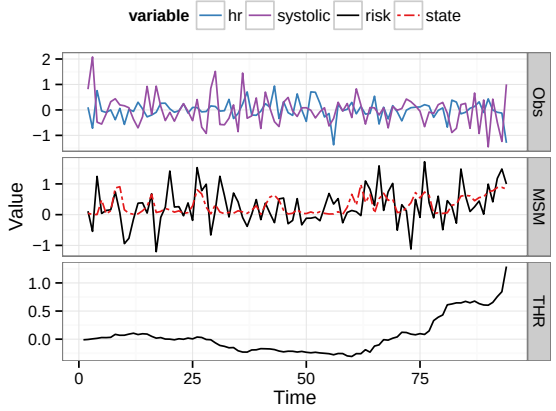

Figure 9: The estimated CAT based on a single patient's observations for both DYNACARE models.

a standard logistic regression model were evaluated on the 20 time periods prior to cardiac arrest. Leave-one-out cross validation was used; each patient trained on the remaining 107 patients. For the "new patient", stratified bootstrapping was used to draw β and then 'unsupervised" particle smoothing was used to estimate the risk trajectory. No cardiac arrest information was provided to the unsupervised particle smoother. Algorithm 4 outlines the general algorithm for estimating CAT for a new patient.

Algorithm 4 DYNACARE estimation algorithm

Find stratum with matching patient age and gender
Draw $\boldsymbol{\beta}^{(i)}$ from stratum of learned parameters
Estimate $r_{1:T}^{(i)}$ using $\boldsymbol{\beta}^{(i)}$
Compute $\hat{r}_{1:T} = E_i[r^{(i)}|\boldsymbol{\beta}^{(i)}]$

4.3 Results

Figure 9 demonstrates the estimated risk trajectory from MSM and THR based on a patient's sequence of observations. The patient-specific model parameters were learned using the general DYNACARE algorithm (Algorithm 1). These parameters were then used to estimate the patient's CAT using an unsupervised particle smoother. Patients generally had similar estimated CAT to the one shown in Figure 9 even when the observations did not exhibit the same temporal patterns.

The distribution of learned model parameters for the

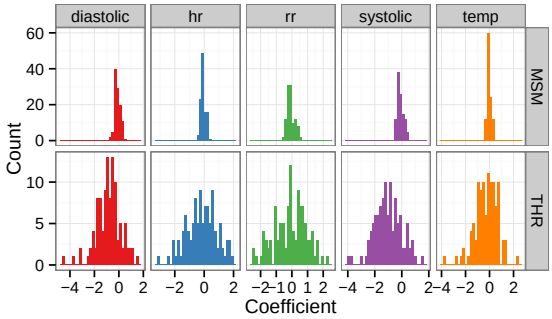

Figure 10: The empirical distribution of parameter values ($\hat{\boldsymbol{\beta}}$) for the DYNACARE models.

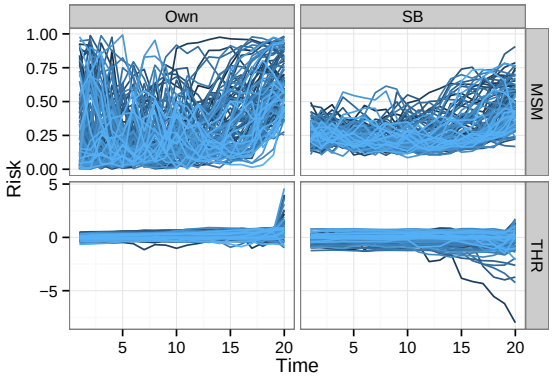

Figure 11: The CAT on the last 20 time periods for the DYNACARE models using either individually learned or stratified bootstrapping β parameters. Cardiac arrest occurs within 4 hours after index 20.

DYNACARE models is shown in Figure 10. Both models have approximately the same mean parameter value for heart rate, respiratory rate, and temperature. MSM parameters have smaller variance. The THR coefficients for blood pressure (diastolic and systolic) are slightly more negative in comparison to their MSM counterparts. This suggests that THR places more weight on the value of these measurements, searching for a downward trend of the blood pressure values.

Differences in the estimated CAT using the "true" learned parameter values and the stratified bootstrapping parameters can be seen in Figure 11. The estimated trajectory for MSM using individually learned parameters is actually more noisy than stratified bootstrapping. However, the opposite occurs for the THR which has less variability for "true" parameter values. The disparity maybe a manifestation of the larger variance in the THR parameter distributions shown in Figure 10.

To create a fair comparison of the DYNACARE models, an unsupervised simple dynamic linear model (DLM), and a standard logistic regression model trained only on the current observations, additional lo-

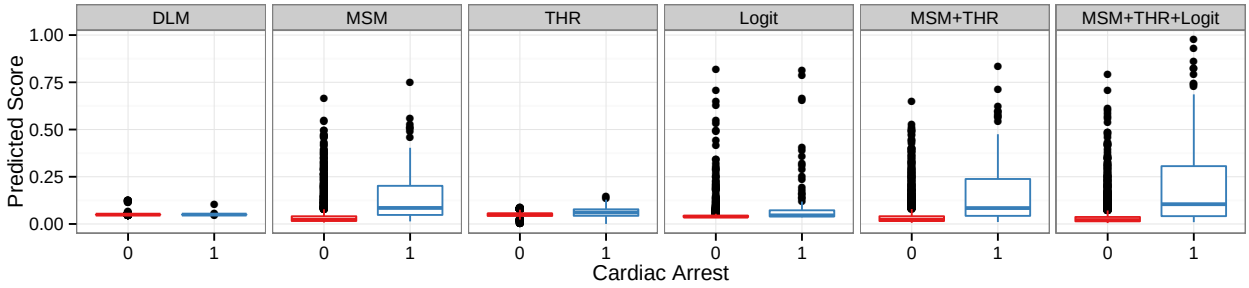

Figure 12: A boxplot of the predictive scores for cardiac arrest. Class 1 represents the time indices right before the cardiac arrest event and class 0 for the remaining time periods. Ensemble scores are computed using a convex combination.

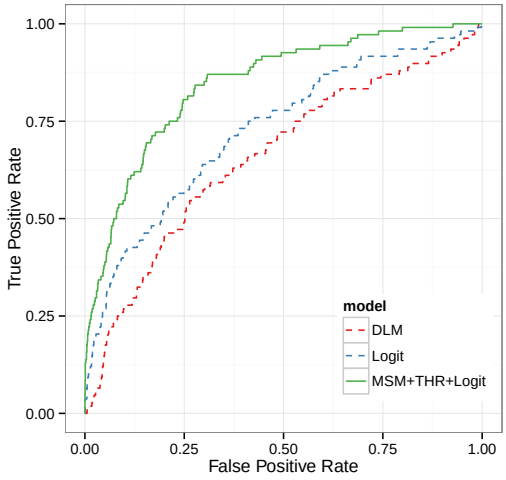

Figure 13: The area under the Receiver Operating Characteristic curve (AUROC) for the simple dynamic linear model, standard logistic model and an ensemble of logistic with the DYNACARE models.

gistic regression models were trained for DLM, MSM, and THR to produce a probability of cardiac arrest based on the risk trajectory value. Figure 12 and Figure 13 illustrate the predictive performance of the models. MSM performs the best of the four models followed by THR. The standard logistic regression model results in a higher number of false positives while the simple DLM performs the worst. Figure 12 also demonstrates an improvement using an ensemble approach of the DYNA-CARE models. However, the ensemble of all three models yields the best predictive performance.

5 Discussion

DYNACARE provides a general methodology for analyzing several types of complex temporal data. The semi-supervised framework allows latent factors to be related to a rare event. Consequently, DYNACARE offers interpretability of the latent factor as well as the predictability of the cardiac arrest event.

The DYNACARE models produce a cardiac arrest trajectory with predictive capability that can be easily visualized and interpreted by a medical professional. The general DYNACARE algorithm allows the model to continuously track a patient's trajectory in real-time using a distributed system. Moreover, the model stores the learned parameters to estimate the risk trajectory for a new patient with limited observations. Furthermore, DYNACARE provides a personalized dynamic hazard function, which cannot be obtained using traditional survival analysis.

This paper introduced two novel dynamic models to estimate a patient's risk of cardiac arrest. The DYNACARE algorithm can be extended to utilize sequential learning of the model parameters. Additionally, DYNACARE models can be augmented to encompass different types of data (categorical or binomial data) and incorporate other features such as additional physiological measurements, laboratory test results, drug dosages, and nurse's notes. Based on the improved performance of the ensemble of MSM, THR, and logistic regression, future work can focus on creating a single model that combines the models simultaneously. Finally, a general framework can be developed to address other maladies (e.g. pneumonia, sepsis, or heart attacks).

In conclusion, we demonstrated the potential of using dynamic models to estimate a patient's risk of cardiac arrest. The results show promise in their ability to accurately identify patients at risk of cardiac arrest, potentially improving the survival rate of ICU patients.

Acknowledgments

This work is supported by a grant from ORNL/CMS and the Schlumberger Centennial Chair in Engineering.

References

A. C. Andréasson, J. Herlitz, A. Bång, L. Ekström, J. Lindqvist, G. Lundström, and S. Holmberg. Characteristics and outcome among patients with a suspected in-hospital cardiac arrest. *Resuscitation*, 39 (1-2):23–31, Oct. 1998.

I. Batal, D. Fradkin, J. Harrison, F. Moerchen, and M. Hauskrecht. Mining recent temporal patterns for event detection in multivariate time series data. In *KDD '12: Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 280–288. ACM, 2012.

L. E. Calvet and A. J. Fisher. How to Forecast Long-Run Volatility: Regime Switching and the Estimation of Multifractal Processes. *Journal of Financial Econometrics*, 2(1):49–83, Jan. 2004.

C. M. Carvalho, H. F. Lopes, and O. Aguilar. Dynamic Stock Selection Strategies: A Structured Factor Model Framework. *Bayesian Statistics*, 9, 2011.

M. M. Churpek, T. C. Yuen, M. T. Huber, S. Y. Park, J. B. Hall, and D. P. Edelson. Predicting Cardiac Arrest on the Wards: A Nested Case-Control Study. *Chest*, 141(5):1170–1176, May 2012.

C. Cox. Threshold dose-response models in toxicology. *Biometrics*, 43(3):511–523, Sept. 1987.

D. R. Cox. Regression models and life-tables. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 187–220, 1972.

A. Doucet and A. M. Johansen. A Tutorial on Particle Filtering and Smoothing: Fifteen years later, Dec. 2008.

A. S. Fialho, F. Cismondi, S. M. Vieira, J. M. C. Sousa, S. R. Reti, and M. D. Howell. Predicting outcomes of septic shock patients using feature selection based on soft computing techniques. In *Information Processing and Management of Uncertainty in Knowledge-Based Systems. Applications*, pages 65–74. Springer Berlin Heidelberg, 2010.

T. J. T. Hodgetts, G. G. Kenward, I. G. I. Vlachonikolis, S. S. Payne, and N. N. Castle. The identification of risk factors for cardiac arrest and formulation of activation criteria to alert a medical emergency team. *Resuscitation*, 54(2):125–131, Aug. 2002.

C. E. Kennedy and J. P. Turley. Time series analysis as input for clinical predictive modeling: modeling cardiac arrest in a pediatric ICU. *Theoretical biology & medical modelling*, 8:40, 2011.

S. Kim, N. Shephard, and S. Chib. Stochastic volatility: likelihood inference and comparison with ARCH models. *The Review of Economic Studies*, 65(3): 361–393, 1998.

J. McBride, D. Knight, J. Piper, and G. B. Smith. Long-term effect of introducing an early warning score on respiratory rate charting on general wards. *Resuscitation*, 65(1):41–44, Apr. 2005.

M. Saeed, M. Villarroel, A. T. Reisner, G. Clifford, L.-W. Lehman, G. Moody, T. Heldt, T. H. Kyaw, B. Moody, and R. G. Mark. Multiparameter Intelligent Monitoring in Intensive Care II: A public-access intensive care unit database. *Critical Care Medicine*, 39(5):952–960, May 2011.

C. Sandroni, G. Ferro, S. Santangelo, F. Tortora, L. Mistura, F. Cavallaro, A. Caricato, and M. Antonelli. In-hospital cardiac arrest: survival depends mainly on the effectiveness of the emergency response. *Resuscitation*, 62(3):291–297, Sept. 2004.

C. Sandroni, J. Nolan, F. Cavallaro, and M. Antonelli. In-hospital cardiac arrest: incidence, prognosis and possible measures to improve survival. *Intensive Care Medicine*, 33(2):237–245, Feb. 2007.

A. F. Smith and J. Wood. Can some in-hospital cardio-respiratory arrests be prevented? A prospective survey. *Resuscitation*, 37(3):133–137, June 1998.

F. Wang, N. Lee, J. Hu, J. Sun, and S. Ebadollahi. Towards heterogeneous temporal clinical event pattern discovery: a convolutional approach. In *KDD '12: Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, 2012.