DYNACARE-OP: Dynamic Cardiac Arrest Risk Estimation Incorporating Ordinal Features

Joyce C. Ho JOYCEHO@UTEXAS.EDU
Yubin Park YUBIN.PARK@UTEXAS.EDU
Carlos M. Carvalho CARLOS.CARVALHO@MCCOMBS.UTEXAS.EDU
Joydeep Ghosh GHOSH@ECE.UTEXAS.EDU
University of Texas at Austin, Austin, TX 78712 USA

Abstract

Cardiac arrest, a deadly condition caused by a sudden failure of the heart, is synonymous with clinical death (in-hospital mortality rate of $\sim 80\%$). Early and accurate estimation of 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 and/or leverage ordinal measurements. This paper presents a dynamic cardiac risk estimation model using ordered probit (DYNACARE-OP) to incorporate ordinal features. The model tracks a patient's risk trajectory, leverages continuous and ordinal clinical measurements, provides an intuitive visualization to medical professionals, improves cardiac arrest event predictability, and estimates the cardiac arrest risk for a new patient.

1. Introduction

Cardiac arrest prevents proper blood circulation due to a sudden failure of the heart function. The heart's pumping action may halt from abnormal rhythms caused by disturbances in the electrical system of the heart. Cardiac arrest is synonymous with death given the mortality rate of $\sim$ 80% (Sandroni et al., 2007). Recent studies have shown that $\sim$ 62% of cardiac arrests could have 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). A quick response to cardiac arrest can also 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, early and accurate identification of at-risk patients is critical to cardiac arrest prevention and improving the survival rate.

Several early warning scores or criteria have been established to predict patients at high risk of experience a cardiac arrest. These systems are designed to detect patient deterioration and alert an emergency response team (Smith & Wood, 1998; Hodgetts et al., 2002; McBride et al., 2005; Churpek et al., 2012). However, scoring systems or activation criteria are unable to capture temporal patterns in clinical measurements. Supplementing the feature set with clinically relevant latent variables, trend and seasonality features (Kennedy & Turley, 2011), or searching for temporal patterns within the data (Batal et al., 2012; Wang et al., 2012) have been suggested as effective ways to incorporate temporal information. Most recently, a dynamic cardiac arrest risk estimation (DYNACARE) model was developed based on stochastic volatility time series models developed for economic forecasting (Ho et al., 2013). Although DYNACARE was able to improve prediction accuracy, the features were restricted to continuous clinical measurements, constraining the feature space and potentially limiting its effectiveness. Moreover, the model cannot be easily adapted for other medical diagnosis that rely primarily on ordinal clinical measurements such as level of consciousness (AVPU scale), pain scores, Modified Early Warning Score (MEWS), and the Sequential Organ Failure Assessment (SOFA) score.

This paper presents a dynamic cardiac arrest risk estimation using the ordered probit (DYNACARE-OP), which substantially extends the scope of the DYNACARE model. The DYNACARE-OP model (i) can

Proceedings of the 30^{th} International Conference on Machine Learning, Atlanta, Georgia, USA, 2013. JMLR: W&CP volume 28.

handle continuous and ordinal features, (ii) improve the tracking of cardiac arrest trajectory, (iii) estimate the risk of two common cardiac arrest categories, (iv) maintain all the existing benefits of the DYNACARE model, and (v) generalize to detecting other medical events.

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-OP Model

2.1. DYNACARE Markov Switching Model (Ho et al., 2013)

The DYNACARE Markov switching model (MSM) represents a patient's cardiac arrest trajectory (CAT) as *a single latent factor* r_t, that governs the gradient of the continuous clinical measurements. The model assumes two heterogenous dynamic structures of CAT, namely a healthy and risky state, where the state influences the mean of the latent factor. Fluctuations around the mean are captured by an error term, or the risk residuals. DYNACARE MSM extends the general dynamic linear model using a stochastic volatility model and a semi-supervised framework to relate the latent factor to the cardiac arrest event.

The use of the stochastic volatility model is motivated by auto-correlation of the risk residuals, where large residuals are clustered together and small residuals follow small residuals. The "volatility clustering" suggests modeling variance using an underlying stochastic process. An alternative to the stochastic process is the use of the generalized autoregressive conditional heteroskedasticity (GARCH) model (Bollerslev, 1986). However, Kim et al. (1998) showed that although the GARCH model has a slightly higher likelihood ratio, SV is a more parsimonious model. In addition, the DYNACARE volatility term not only models the autocorrelation amongst risk factor residuals (time to time) but captures individual differences in the risk residuals (patient to patient). Although the volatility term introduces non-linearity into the system, a particle filter (a sequential Monte Carlo sampling method) can be used to estimate the latent variables (Doucet & Johansen, 2008).

The semi-supervised framework uses partial knowledge to connect a patient's cardiac arrest risk to the latent factor. The only provided information is the fact that a cardiac arrest event occurred at a specific time point. During the remaining observation period, we cannot ascertain the period in which the patient was healthy or if any unrecorded or unobserved cardiac arrest events transpired. Thus, DYNACARE MSM sets the state to be risky during the single known cardiac event. To utilize this information, "backward smoothing" needs to be performed to infer the latent factor using all the observations including "future" measurements. A particle filter that incorporates the forward-backward recursion is known as particle smoothing. This semi-supervised framework allows interpretability and predicability of a cardiac arrest event.

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

Equation block 1 summarizes the DYNACARE MSM model formulation. DYNACARE MSM can estimate a patient's cardiac arrest trajectory, model a new patient's trajectory with a limited number of observations, and predict a cardiac arrest event time. However, the model only allows the observations to be numeric responses, placing an unnecessary restriction on the feature space.

2.2. DYNACARE-OP

Ordinal measurements can be introduced into the existing DYNACARE model using several techniques. The simplest approach is to convert the ordinal measurements to integer values. While no modification is required, the sequence of ordinal measurements, or a set of step responses, cannot be easily approximated by the linear observation equation. Another method is to use a different $\boldsymbol{\beta}$ for each ordinal response that is shared across all patients. However, a unique $\boldsymbol{\beta}$ needs to be learned for each possible combination of ordinal measurements, leading to "curse of dimensionality" related problems. Our model uses a data augmentation and ordered probit trick (Albert & Chib, 1993) to incorporate ordinal measurements.

DYNACARE-OP partitions the set of observations $\mathbf{y}$ into continuous measurements $\mathbf{y^c}$ and ordinal measurements $\mathbf{y^o}$. A new latent continuous variable $\mathbf{z}$ is introduced to link a linear regression to the ordinal response outcomes, resulting in easy-to-implement simulations from standard distributions. Suppose that a single ordinal feature y_ρ has J ordered categories.

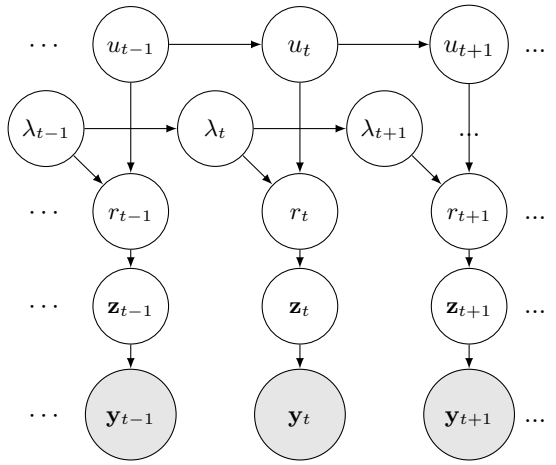

Figure 1. The graphical representation of the DYNACARE-OP model.

We observe $y_\rho = j$ when the latent variable z_ρ falls within the range $\gamma_{j-1} < z_\rho \leq \gamma_j$, where z_ρ is normally distributed around $\beta_\rho r$. To ensure identifiability, we define $\gamma_0 = -\infty$ and $\gamma_J = \infty$. For the continuous features, $\mathbf{z^c}$ is simply the observation $\mathbf{y^c}$.

In addition to the new latent variable $\mathbf{z}$, DYNACARE-OP assumes the healthy and risky states govern the observations directly to avoid modeling gradients of ordinal measurements $\mathbf{\Delta y^o}$. Continuous measurements are standardized per patient such that $\mathbf{y^c}$ has a mean of 0 and variance of 1, minimizing the effect of patient offsets and potential scaling issues. The model is illustrated in equation block 2, where λ is the stochastic volatility term, u is the latent state, r is the latent factor CAT, $\mathbf{z}$ is the data augmentation variable, and$\mathbf{y}$ is the set of observations that can be decomposed into the continuous, $\mathbf{y^c}$, and ordered categorical, $\mathbf{y^o}$, measurements. Figure 1 shows the graphical representation of the model.

$$
\begin{aligned}
\lambda_t &= \lambda_{t-1} + \delta_t & \delta_t &\sim N(0, k^2) \\
u_t &\sim \text{MarkovChain}(u \mid u_{t-1}) & u_t &\in \{s_h, s_c\} \\
r_t &= \alpha_{u_t} + \varepsilon_t & \varepsilon_t &\sim N(0, \exp(\lambda_t)) \\
\alpha_{u_t} &\in \{\alpha_{s_h}, \alpha_{s_c}\},\ \alpha_{s_h} \neq \alpha_{s_c} & & \\
\mathbf{z}_t &= \boldsymbol{\beta} r_t + \boldsymbol{\eta}_t & \boldsymbol{\eta}_t &\sim N(0, \boldsymbol{\Sigma}) \\
\boldsymbol{\Sigma} &= \text{diag}(\sigma_1^2, \cdots, \sigma_f^2) & & \\
\mathbf{y}_t &= \begin{bmatrix} \mathbf{y}_t^c & \mathbf{y}_t^o \end{bmatrix}^T & & \\
\mathbf{y}_t^c &= \mathbf{z}_t^c & & \\
\mathbf{y}_t^o &= o_j, \quad \gamma_{j-1}^o < z_t^o \leq \gamma_j^o & &
\end{aligned}
\tag{2}
$$

The joint distribution of DYNACARE-OP is shown in Table 1. DYNACARE-OP uses the expectation maximization algorithm to estimate the parameters $\boldsymbol{\beta}, \boldsymbol{\Sigma}$.

$$
\begin{aligned}
\{\hat{r}_{1:T}, \hat{u}_{1:T}, \hat{\lambda}_{1:T}\} &\sim E[r_{1:T}, u_{1:T}, \lambda_{1:T} | \mathbf{z}_{1:T}, \boldsymbol{\beta}, \boldsymbol{\Sigma}] \\
\{\hat{z}_{1:T}\} &\sim E[z_{1:T} | \mathbf{y}_{1:T}, r_{1:T}, \boldsymbol{\beta}, \boldsymbol{\Sigma}, \gamma^{\mathbf{o}}] \\
\{\boldsymbol{\beta}, \boldsymbol{\Sigma}\} &\sim \max p(\mathbf{z}_{1:T}, r_{1:T}, u_{1:T}, \lambda_{1:T} | \boldsymbol{\beta}, \boldsymbol{\Sigma})
\end{aligned}
\tag{3}
$$

The latent variables r, u (equation 3) can be efficiently simulated using a particle smoother. The forward and backward smoothing equations are summarized in Table 1. The DYNACARE-OP semi-supervised learning incorporates two facts: (i) a cardiac arrest event occurs during the last time period and (ii) a patient exhibits clinical deterioration at least 4 hours before arrest. Thus, the latent state u is set to the risky state during the time blocks 4 hours before cardiac arrest (equation 4). The semi-supervised framework can also be used as a knob for controlling the cardiac arrest notification time. For earlier prediction, u should be set to the risky state for the desired detection time.

$$
p(u_t = s_c | \mathbf{y}_t) \approx 1, \quad \forall t > t_{\text{CA}} - 4 \tag{4}
$$

To prevent particle degeneracy introduced by the backward smoothing process for $t \ll t_{\text{CA}}$, where a single unique particle approximates the latent variable distribution, DYNACARE-OP uses the simplest approach, a fixed-lag approximation (Doucet & Johansen, 2008). This method leverages the forgetting properties of hidden Markov models such that for a large enough Δ, $p(r_{1:n} | \mathbf{y}_{1:T}) \approx p(r_{1:n} | \mathbf{y}_{\min(n+\Delta, 1:T)})$. Algorithm 1 outlines the particle smoother.

Algorithm 1 DYNACARE-OP particle smoother

Draw $k^{(i)} \sim \Gamma^{-1}(\alpha, \beta)$
$\lambda_0^{(i)} \sim \mathcal{N}(0, k^{(i)})$, $u_0^{(i)} \in \{0, 1\}$, $r_0^{(i)} \sim \mathcal{N}(0, 1)$
for $t = 1 : T$ **do**
  for $\tau = t : \min(t + L, T)$ **do**
    Draw $\lambda_\tau^{(i)} \sim \mathcal{N}(\lambda_{\tau-1}^{(i)}, k^{(i)})$
    Draw $u_\tau^{(i)} \sim \text{MarkovChain}(u \mid u_{\tau-1}^{(i)})$
    Calculate $r_{\min}$ and $r_{\max}$ using $\mathbf{y}_\tau^{\mathbf{o}}, \boldsymbol{\gamma}^o, \boldsymbol{\beta}^o$
    Draw $r_\tau^{(i)} \sim \mathcal{TN}(\alpha_{u_\tau^{(i)}}, \exp(\lambda_\tau^{(i)}/2), r_{\min}, r_{\max})$
    $w_\tau^{(i)} \propto \exp(\frac{1}{2}(\mathbf{y}_\tau - \boldsymbol{\beta} r_\tau^{(i)})^\top \boldsymbol{\Sigma}^{-1}(\mathbf{y}_\tau - \boldsymbol{\beta} r_\tau^{(i)}))$
  end for
  for $\tau = \min(t + L, T) : t$ **do**
    Compute $w_{\tau-1}^{(i)}$ according to Table 1
  end for
  $\hat{u}_t = \sum w_t^{(i)} u_t^{(i)}$ and $\hat{r}_t = \sum w_t^{(i)} r_t^{(i)}$
end for

The DYNACARE-OP particle smoother calculates the upper and lower bounds on the latent variable r given the ordinal observations y^o, the bin boundaries boundaries $\boldsymbol{\gamma}^o$, and the parameters $\boldsymbol{\beta}^o$ and draws from a

Table 1. Important DYNACARE-OP distributions

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

truncated normal distribution. To minimize repeated computation of the range introduced by the backward smoothing step ($\mathcal{O}(L)$), DYNACARE-OP computes $r_{\min}$ and $r_{\max}$ for all the steps during initialization, incurring an additional memory cost of $\mathcal{O}(T)$ per patient.

Patient-specific individual bin boundaries, $\boldsymbol{\gamma^o}$, are impractical for medical settings. A patient may not have observations that span the J categories for each ordinal measurement type. Consequently, a normalization factor for the $\boldsymbol{\beta}^o$ needs to be computed to allow for parameter comparison. Additionally, learned model parameters cannot be used to estimate the risk of a new patient unless the model knows apriori the observed responses. Thus, DYNACARE-OP defines the $\boldsymbol{\gamma}^o$ boundaries using the empirical distribution of all patient measurements. Patient-specific $\boldsymbol{\beta}^o$ then learned according to the shared global bin boundaries.

DYNACARE-OP combines expectation maximization (EM) sampling and particle smoothing to estimate a patient's cardiac arrest risk. For each patient, the particle smoother with fixed-lag approximation is used to estimate the latent variables r and u. The data augmentation variable z is then updated using $\boldsymbol{\beta}$ and r. Model parameters are then estimated from all the latent variables. The process iterates until convergence of the parameters occurs. The DYNACARE-OP framework is outlined in Algorithm 2.

Algorithm 2 DYNACARE-OP algorithm

```
Calculate γ^o using empirical distributions
for patient=1 : N do
  while β_i not converged do
    Estimate û_{1:T}, r̂_{1:T} with Algorithm 1
    Estimate ẑ_t^o ~ TN(β^o r̂_t, Σ^o, γ^o_{j-1}, γ^o_j)
    Learn β_i, Σ_i given r̂_{1:T}, ẑ_{1:T}
  end while
end for
```

Learned patient parameters are stored in a database to model new patients. For a new patient, stratified bootstrapping is used to draw model parameters $\boldsymbol{\beta}, \boldsymbol{\Sigma}$ based on age and gender. An "unsupervised" particle smoother is used to estimate the latent states, which does not incorporate the cardiac arrest information (equation 4). The new patient's cardiac arrest trajectory is the average of the estimated states from the bootstrap samples. Algorithm 3 illustrates the procedure for estimating the cardiac arrest trajectory for a new patient.

Algorithm 3 DYNACARE-OP estimation algorithm

```
Draw β^(i), Σ^(i) from stratum of learned parameters
Estimate r_{1:T}^(i) using β^(i), Σ^(i)
Compute r̂_{1:T} = E_i[r_{1:T}^(i) | β^(i), Σ^(i)]
```

DYNACARE-OP retains all the benefits of DYNACARE. Since the particle smoother (Algorithm 1) and estimation algorithm (Algorithm 2) is performed on a patient-level, the computation can be distributed across multiple machines. DYNACARE-OP also provides a personalized dynamic hazard function that varies over time and patient history using a derivation similar to the DYNACARE MSM model. The DYNACARE-OP hazard function, or the probability of a cardiac arrest event at time t, is the transition probability from $\hat{u}_t$ to the risky state given all the observations up to the current time period $y_{1:t}$. Therefore the personalized hazard function is $h(t) = p(\hat{u}_t = s_c|y_{1:t})$.

Thus DYNACARE-OP can (i) learn an individual patient's model parameters, (ii) estimate the cardiac arrest trajectory, (iii) model a new patient with limited number of observations, (iv) deliver instantaneous results for a large patient population via distributed computing, and (v) provide a personalized dynamic hazard function.

3. Empirical Studies

3.1. Data

The Multiparameter Intelligent Monitoring in Intensive Care II (MIMIC-II) database, developed to sup-

Table 2. Clinical Measurements

Name	Type
Pulse pressure (pp)	Continuous
Pulse pressure index (ppi)	Continuous
Respiratory rate (rr)	Continuous
Heart rate (hr)	Continuous
Temperature (temp)	Continuous
Glasgow Coma Score (gcs)	Ordinal
Respiratory pattern (resp pattern)	Binary
Heart rhythm (heart rhythm)	Binary

port research in clinical decision support and critical care medicine (Saeed et al., 2011), is the most extensive and publicly available intensive care unit (ICU) resource. The database contains over 30,000 ICU patients during a 6-year period from Boston's Beth Israel Deaconess Medical Center, providing a good evaluation for our model.

The study was conducted on patients at least 18 years of age at admission time who experienced either an asystole or ventricular tachycardia (ventachy) cardiac arrest event. Asystole, a sudden pause of heart muscle contractions, and ventricular tachycardia, an irregular heartbeat caused by a fast heart rate are two common cardiac arrest categories. We focused on eight clinical measurements common in these patients which are summarized in Table 2. Respiratory pattern and heart rhythm were grouped into binary outcomes, normal and abnormal, and treated as ordinal measurements with 2 response types. Data was discretized into 2-hour bins prior to the first observed cardiac arrest event, with the event denoted as time 0. Additionally, we required each patient to have at least 72 hours of measurements prior to the cardiac arrest event to ensure sufficient data points.

In the MIMIC-II database, 725 cardiac arrest patients were diagnosed with asystole or ventricular tachycardia. 298 of these patients met the minimum data requirements where 162 patients experienced an asystole event during their ICU stay. On average, patients had 160 hours of data (∼1 week) with a standard deviation of 48 hours. We assumed unobserved measurements signified no change in the status and used the zero-order hold (Fialho et al., 2010), maintaining the last observed value. Figure 2 shows a plot of the last 200 hours prior to the cardiac arrest event for one of the patients.

3.2. Results

The distribution of learned parameters, $\boldsymbol{\beta}$, for the cardiac arrest categories is shown in Figure 3. Both cardiac arrest types seem to have similar dynamics except

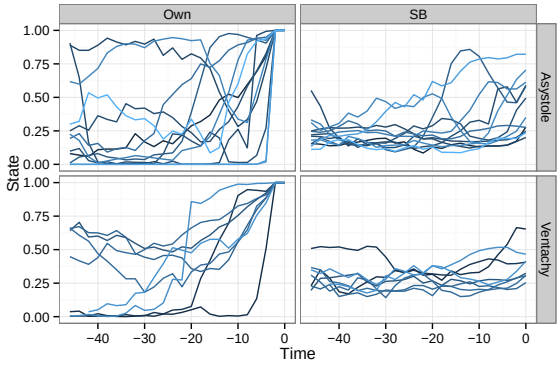

Figure 5. 20 samples of the estimated state u on the last 48 hours prior to the cardiac arrest event using individually learned (own) or stratified bootstrapping (SB) β parameters based on the observations.

for pulse pressure, heart rate, Glasgow Coma Score (GCS), and heart rhythm. The pulse pressure parameters are lower for asystole cardiac arrest patients. The GCS parameters also tend to be more positive in asystole patients, suggesting a higher cardiac arrest risk for comatose patients. Prior research has indicated that GCS is a strong outcome assessor for cardiac arrest survivors (Schefold et al., 2009), but the reverse which is suggested by our models, has not been shown before. For ventricular tachycardia patients, the heart rate parameters are slightly skewed positively while the heart rhythm parameters have two modes far away from 0. The figure also implies that pulse pressure index, respiratory rate, and temperature are not important features in estimating a patient's cardiac arrest trajectory with a distribution centered around 0.

Differences in the estimated state using the individually learned, or "true", and the stratified bootstrapping β parameters are illustrated in Figure 5. The estimated state using the "true" parameters is noisier than the stratified bootstrapping, and tends to have a sudden shift to the risky state near the arrest time (time 0). The figure illustrates a "smoothing" effect from other patients for the stratified bootstrap with a gradual rise in the estimated state. In addition, the stratified bootstrapping states for ventricular tachycardia patients do not have a noticeable separation in the cardiac arrest trajectory close to the cardiac arrest time. Figure 5 suggests that detecting ventricular tachycardia may be harder and requires additional features.

A comparison of the DYNACARE MSM model and the DYNACARE-OP model is shown in Figure 4. The contour plots demonstrate the benefits of incorporating ordinal features to estimate the cardiac arrest.

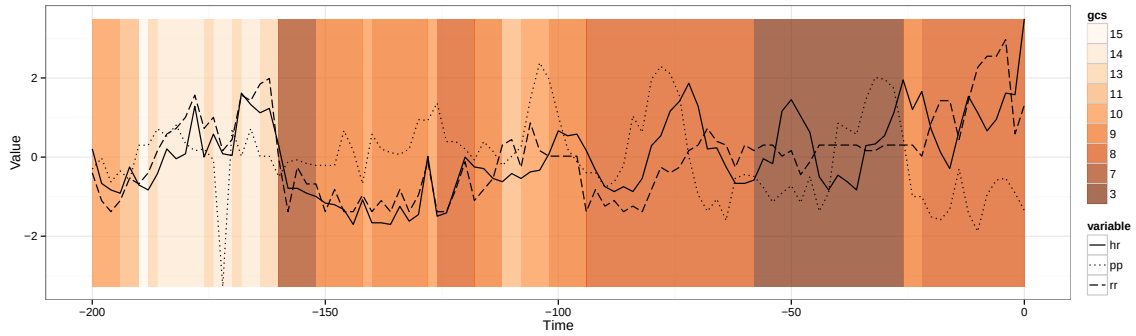

Figure 2. An example of a patient's normalized measurements for heart rate, pulse pressure, respiratory rate and the Glasgow Coma Score values.

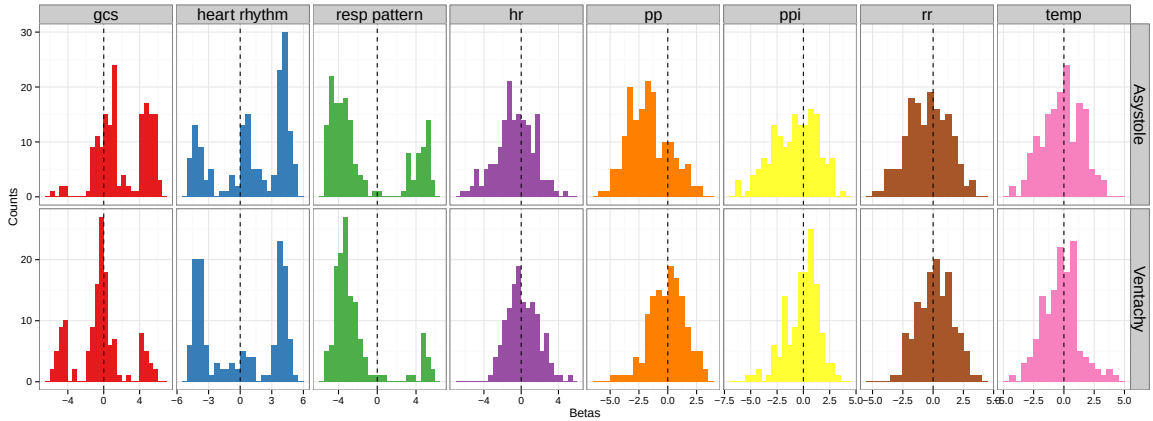

Figure 3. The distribution of β parameters for each feature based on the cardiac arrest type.

The additional constraints on the risk trajectory introduced by the ordinal observations helps the model better estimate the patient's risk. Although DYNACARE MSM estimates larger risk values for the "more certain" ($u > 0.60$) patients close to the cardiac arrest event, the DYNACARE-OP states generally exceed the 0.5 threshold. However, DYNACARE-OP generally estimates higher risk at the lower, more "uncertain" regimes, potentially boosting the detection rate. It is also important to note that the density estimates 36 hours prior to the event ($t = -36$) show that the DYNACARE-OP model does not arbitrarily boost the estimated state as the points are distributed around the dotted line.

The different $\boldsymbol{\beta}$ parameter distributions in Figure 3 suggest the potential to estimate the cardiac arrest risk for each category. Figure 6 shows the kernel density estimation for the last 4 hours using stratified bootstrapping with the same cardiac type (e.g. asystole patients with asystole parameters) and the other cardiac type (e.g. asystole patients with ventricular tachycardia parameters). For the more "uncertain" regimes, asystole cardiac arrest can also be estimated using the ventricular tachycardia parameters. However, ventricular tachycardia patients tend to have a lower estimated risk when using the asystole patient parameters.

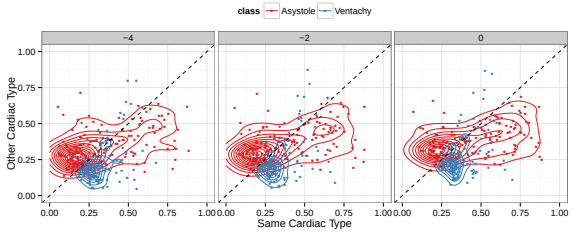

Figure 6. A kernel density estimation of the patient's state using stratified bootstrapping parameters using the same cardiac arrest category or the other cardiac arrest type.

The predictive performance of the DYNACARE-OP model was evaluated against a standard logistic regression and linear kernel support vector machine trained only on the current observations and the DYNACARE MSM model. To obtain a probability of cardiac arrest for the DYNACARE-OP and DYNACARE MSM model, a logistic regression model was trained only

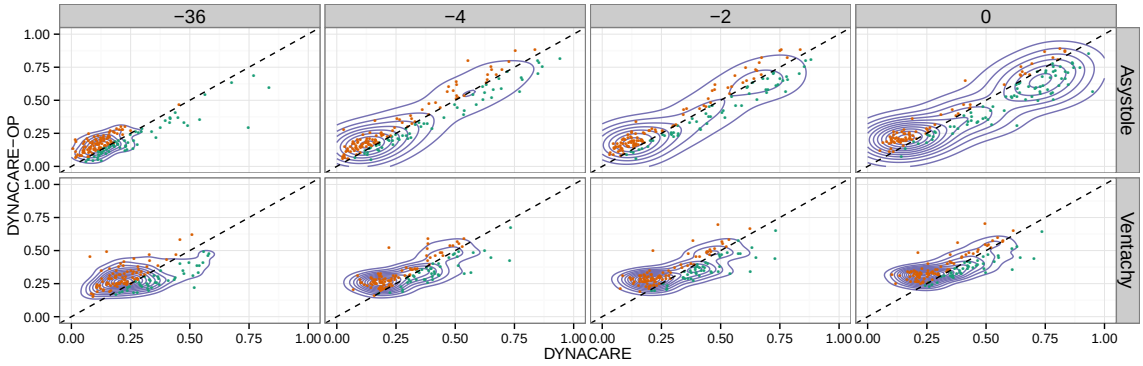

Figure 4. The contours from the kernel density estimation of the patient's state using stratified bootstrapping parameters for the DYNACARE MSM and the DYNACARE-OP models. Orange points denote patients whose state has increased under the DYNACARE-OP, while green points represent a higher estimated DYNACARE state. The dotted line denotes the location where both models have the same estimated state value.

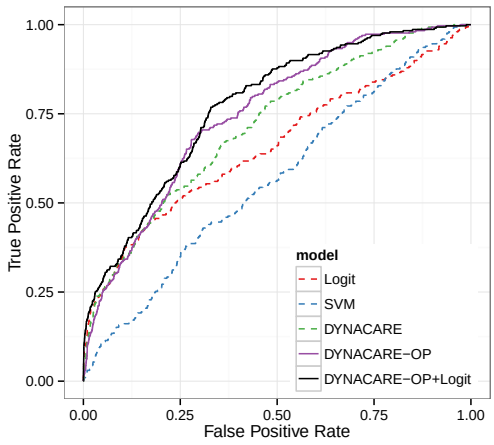

Figure 7. The area under the Receiver Operating Characteristic curve (AUROC) for the standard logistic model (Logit), support vector machine (SVM), DYNACARE MSM model, DYNACARE-OP model, and an ensemble of logistic with the DYNACARE-OP model.

on the estimated state value. Figure 7 demonstrates the predictive performance of the 4 models in addition to an ensemble of the logistic regression and DYNACARE-OP model. Generally, DYNACARE-Op outperforms the other three models except in the low false positive regime. Furthermore, the ensemble of the standard logistic regression with just the current features and the DYNACARE-OP estimated state yields the best overall predictive performance.

4. Discussion

The DYNACARE-OP model provides a formulation for incorporating ordinal features to dynamically estimate an individual patient's cardiac arrest risk. The model can not only utilize continuous and ordinal clinical measurements but also estimate the risk trajectory for a new patient without requiring prior knowledge of the ordinal measurement response types. Moreover, the DYNACARE-OP framework can be easily adapted to track the risk for other catastrophic medical events, such as septic shock or mortality. Other diseases also have risk scores that are actually ordinal, for example a variety of self-reported pain scores, the APGAR score for neo-natal health, and the PIM2 score for rating severity of medical illness in children. We expect that DYNACARE-OP will be able to exploit such ordinal score information while building predictive models for the corresponding illnesses. The semi-supervised framework allows us to relate the latent factor to the medical event of interest.

DYNACARE-OP maintains all the benefits of DYNACARE and significantly improves the predictability of a cardiac arrest event. Our model produces a cardiac arrest trajectory that can be easily visualized and interpreted by a medical professional, simultaneously track multiple patients' risk in real-time using a distributed implementation, and provides a personalized hazard function that incorporates observations from ordinal features. Furthermore, DYNACARE-OP can estimate the risk of either an asystole or ventricular tachycardia cardiac arrest event.

DYNACARE-OP can be further extended to encompass additional physiological measurement types, laboratory test results, and clinician's notes. Our model can also be augmented to include interventions such as drugs, surgery, or other "major" events.

In conclusion, we demonstrated a dynamic model that utilizes ordinal features to estimate a patient's risk of cardiac arrest. The results show the benefits of incor-

porating additional features to accurately identify patients at risk of cardiac arrest; potentially preventing unnecessary cardiac arrests and improving the survival rate of ICU patients.

References

Albert, James H and Chib, Siddhartha. Bayesian Analysis of Binary and Polychotomous Response Data. *Journal of the American Statistical Association*, 88(422):669–679, June 1993.

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

Batal, I, Fradkin, D, Harrison, J, Moerchen, F, and Hauskrecht, M. 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*, pp. 280–288. ACM, 2012.

Bollerslev, Tim. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3):307–327, 1986.

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

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

Fialho, Andre S, Cismondi, Federico, Vieira, Susana M., Sousa, Joao M. C., Reti, Shane R., and Howell, Michael D. 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*, pp. 65–74. Springer Berlin Heidelberg, 2010.

Ho, Joyce C, Park, Yubin, Carvalho, Carlos M, and Ghosh, Joydeep. DYNA-CARE: Dynamic Cardiac Arrest Risk Estimation. In *AISTATS'13: Proceedings of the International Conference on Artificial Intelligence and Statistics*, 2013.

Hodgetts, Timothy J TJ, Kenward, Gary G, Vlachonikolis, Ioannis G IG, Payne, Susan S, and Castle, Nicolas N. The identification of risk factors for cardiac arrest and formulation of activation criteria to alert a medical emergency team. *Resuscitation*, 54(2):125–131, August 2002.

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

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

McBride, Jackie, Knight, Debbie, Piper, Jo, and Smith, Gary B. Long-term effect of introducing an early warning score on respiratory rate charting on general wards. *Resuscitation*, 65(1):41–44, April 2005.

Saeed, Mohammed, Villarroel, Mauricio, Reisner, Andrew T, Clifford, Gari, Lehman, Li-Wei, Moody, George, Heldt, Thomas, Kyaw, Tin H, Moody, Benjamin, and Mark, Roger G. Multiparameter Intelligent Monitoring in Intensive Care II: A public-access intensive care unit database. *Critical Care Medicine*, 39(5):952–960, May 2011.

Sandroni, Claudio, Ferro, Giorgia, Santangelo, Susanna, Tortora, Francesca, Mistura, Lorenza, Cavallaro, Fabio, Caricato, Anselmo, and Antonelli, Massimo. In-hospital cardiac arrest: survival depends mainly on the effectiveness of the emergency response. *Resuscitation*, 62(3):291–297, September 2004.

Sandroni, Claudio, Nolan, Jerry, Cavallaro, Fabio, and Antonelli, Massimo. In-hospital cardiac arrest: incidence, prognosis and possible measures to improve survival. *Intensive Care Medicine*, 33(2):237–245, February 2007.

Schefold, Joerg C, Storm, Christian, Krüger, Anne, Ploner, Christoph J, and Hasper, Dietrich. The Glasgow coma score is a predictor of good outcome in cardiac arrest patients treated with therapeutic hypothermia. *Resuscitation*, 80(6):658–661, June 2009.

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

Wang, F, Lee, N, Hu, J, Sun, J, and Ebadollahi, S. 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.