Learning from Different Perspectives: Robust Cardiac Arrest Prediction via Temporal Transfer Learning

Joyce C. Ho[1] and Yubin Park[2]

Abstract— **Predicting and preventing cardiac arrest is one of the biggest challenges of contemporary cardiology, as a patients survival depends on the effectiveness of the emergency response teams. While "black-box models have shown to have better predictive accuracies for cardiac risk stratification, early warning scoring systems are more prominent in the hospital setting due to their ease of implementation and interpretability. We propose a temporal transfer learning approach to utilize information from adjacent time points to yield an early cardiac arrest prediction model that is robust in predictive accuracies as well as maintains the interpretability of the model coefficients. Our model estimates the logistic regression coefficients simultaneously for various time points to share knowledge from different observation windows. This framework can overcome small sample size issues, and result in robust estimation of the model coefficients. We find that our model consistently outperforms a logistic regression model fit only on vital sign data from a single time slice for 763 intensive care unit patients. Moreover, we find that the estimated coefficients from our model captures temporal trends in the data.**

I. INTRODUCTION

One of the great challenges of contemporary cardiology is the prediction and prevention of cardiac arrest. Sudden cardiac death claims approximately 300,000 lives in the United States annually [1], with an in–hospital mortality rate of $\sim 80\%$. Survival depends heavily on the effectiveness of emergency response with the most important intervention is early defibrillation as a patient's survival decreases by 10% with each minute of delay [1], [2]. Cardiac arrest prediction and prevention is an active area of research with an increased focus on risk stratification algorithms as medical resources (e.g., doctors, nurses, beds) are limited.

Current risk–stratification systems are based on clinical judgement and traditional vital signs including heart rate, respiratory rate, blood pressure, and temperature [3], [4], [5]. These early warning systems, while prevalent in many hospital systems due to the ease of implementation, fail to capture temporal patterns in the physiological measurements and also suffer from the ability to accurately identify high-risk patients with sufficient intervention time. Studies have shown that temporal changes can help better characterize changes in the risk of cardiac arrest patients [6], [7], [8], [9]. However, these models are less interpretable in comparison with the risk–stratification systems due to the "black"-box nature. We propose TTL-Reg, our temporal transfer learning based model, which utilizes information from other time points to improve early cardiac arrest prediction while maintaining model interpretability. Our results on 763 intensive care unit (ICU) patients illustrate the benefit of learning from different perspectives to yield a robust early prediction model.

II. TTL-REG

Our objective is to build a cardiac arrest risk prediction model capable of early notification at time z ($z \geq 5$ hours prior to the event). If only data observed at least z hours before the event is used to fit a logistic regression model, not only will the number of samples be relatively low but also result in an overfit model. Figure 1 shows the estimated parameters from six different logistic regression models learned only on data observed at least z hours before the event amongst 763 ICU patients. The estimated coefficients vary significantly from one time to the next even for a relatively small time interval (1 hour). The question then is whether or not information from adjacent time points can be used to improve the coefficient estimation process as well as yield a robust predictive model.

We postulate that knowledge sharing from other time points both directly preceding and succeeding z can be utilized for robust parameter estimation and improve early cardiac arrest prediction. Transfer learning is an area of research in machine learning that focuses on knowledge transfer between different but related problems [10]. Under the transfer learning paradigm, we can pose the estimation of the logistic regression coefficients at the various time points as related problems, which allows the model at time z to learn from different time perspectives. Our model, TTL-Reg, uses a temporal regularization approach to smooth the estimated coefficients between adjacent time points. Thus the coefficient estimation at time z borrows knowledge from the time point before $(z-1)$ and after $(z+1)$.

We have a stream of data $\mathcal{D}$ that contains multiple measurements for each patient and time point $[\mathbf{X}_1, \mathbf{X}_2, \cdots, \mathbf{X}_T]$, where the time point refers to the data that is observed at least t hours before the event time. It is important to note that for the same patient and two different time points, t and $t+1$, the observations $\mathbf{x}_t$ and $\mathbf{x}_{t+1}$ can be drastically different. We denote the loss function of the logistic regression model for a particular time point t as $\ell(\mathbf{y}_t, \mathbf{X}_t \beta_t)$, where $\mathbf{y}_t$ is a binary vector that represents whether or not the patient experiences cardiac arrest t hours later, and β_t are the estimated coefficients for our model. To utilize information

[1]J.C. Ho is the corresponding author for this work. She is with the Department of Mathematics and Computer Science, Emory University, Atlanta, GA 30322, USA. Email: joyce.c.ho@emory.edu

[2]Y. Park is with Accordion Health Inc., Austin, TX 78756, USA. Email: yubin@accordionhealth.com

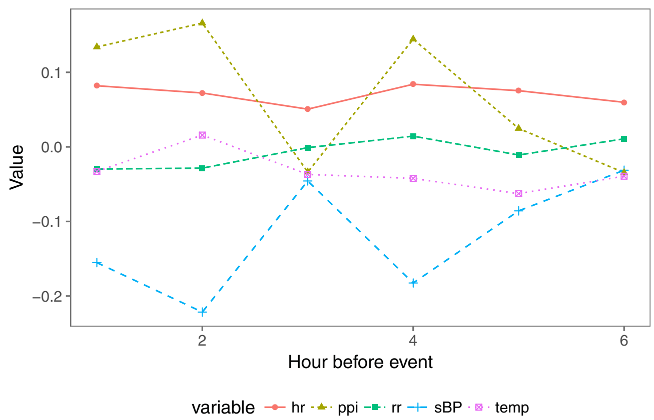

Fig. 1. Trajectories of the estimated coefficients for different time points for one train–test split.

from different time perspectives, we regularize the difference between the coefficients of two adjacent time points, $\lambda||\beta_t - \beta_{t+1}||_2^2$. The regularization parameter, λ, controls the amount of information that is shared between two time points. Higher values of λ will drive the coefficients to be similar to one another, while lower values of λ allow for some freedom between two adjacent time points. Our method for achieving temporal smoothing between the coefficients is related to a technique known as isotropic total variation in the digital image processing world. The objective function for our temporal transfer learning approach is summarized in Equation (1).

$$f(\boldsymbol{\beta}) = \sum_{k=1}^{T} \ell(y_k, \mathbf{X}_k \beta_k) - \sum_{k=1}^{T-1} \frac{\lambda}{2}||\beta_t - \beta_{t+1}||_2^2 - \sum_{k=2}^{T} \frac{\lambda}{2}||\beta_t - \beta_{t-1}||_2^2 \quad (1)$$

As there is no closed form solution for maximizing the logistic loss function for a single time point, let alone multiple time points simultaneously, we take an iterative approach. We solve for each time point, t given the estimated coefficients for the adjacent time points, $t-1$ and $t+1$. Thus at each step, we maximize the function shown in Equation (2) with respect to β_t, with estimates of the other coefficients $\hat{\beta}_{t-1}$ and $\hat{\beta}_{t+1}$.

$$f(\beta_t) = \ell(y_t, \mathbf{X}_t \beta_t) - \frac{\lambda}{2}||\beta_t - \hat{\beta}_{t+1}||_2^2 - \frac{\lambda}{2}||\beta_t - \hat{\beta}_{t-1}||_2^2 \quad (2)$$

We can transform the objective in Equation (2) into a familiar form by introducing a new variable $\alpha = \beta_t - \frac{\beta_{t-1}+\beta_{t+1}}{2}$. The transformed function, illustrated in Equation (3), is simply an ℓ_2-regularized logistic regression model with an offset. Therefore, we can utilize existing software packages (e.g., `glmnet` in R) to estimate the parameters of our model.

$$f(\alpha) = \ell\left(y_t, \mathbf{X}_t \alpha + \mathbf{X}_t(\frac{\beta_{t-1} + \beta_{t+1}}{2})\right) - \lambda||\alpha||_2^2 \quad (3)$$

Algorithm 1: Pseudocode for TTL-Reg

Data: $\mathbf{X}_1, \cdots, \mathbf{X}_T, \mathbf{y}_t, \cdots, \mathbf{y}_T, \lambda$
Result: $\beta_1, \beta_2, \cdots, \beta_T$
while *not converged* ($||\boldsymbol{\beta}^{(i)} - \boldsymbol{\beta}^{(i-1)}||_F$) **do**
  for $t = 1, \cdots, T$ **do**
    `#Calculate offset`
    if $t == 1$ **then**
      offsetAmt = β_{t+1}
    else if $t == T$ **then**
      offsetAmt = β_{t-1}
    else
      offsetAmt = $\frac{\beta_{t-1}+\beta_{t+1}}{2}$
    end
    `# Fit ridge logistic regression`
    logModel = ridge($\mathbf{X}_t$, $\mathbf{y}_t$, λ, offset=-offsetAmt)
    β_t = coef(logModel) + offsetAmt
  end
end

Note that for the corner cases ($t = 1$ or $t = T$), the new variable α is simply defined as the difference between the relevant adjacent time point (i.e., $\alpha = \beta_1 - \beta_2$). The pseudo–code for TTL-Reg is shown in Algorithm 1.

III. EXPERIMENT

A. Data

Our study was conducted on ICU patients who were between the ages of 50 and 75 at the time of admission from the MIMIC-II database [11]. We analyze the following sets of patients:

1) Cardiac arrest patients who had either an asystole or ventricular tachycardia event. The index time for these patients is the time of the first event that is recorded in the event table.
2) Non-cardiac arrest patients who did not experience either an asystole or ventricular tachycardia event. The index time for these patients is randomly sampled from their hospital stay.

We focused on six commonly observed clinical measurements and one derived measurement prior to the index time. Our variables include temperature (temp), peripheral capillary oxygen saturation (spo2), heart rate (hr), respiratory rate (rr), diastolic blood pressure (dBP), systolic blood pressure (sBP), and pulse pressure index (ppi) which is the difference between systolic and diastolic blood pressure over the systolic pressure. We identified 6 time points of interest for our study, which ranged from 1 to 6 hours before the event in hour increments. For each time point, we used the last observed measurement within a 12–hour sliding window to construct our feature matrices $\mathbf{X}_t$. For example, for the time point of 1 hour, we consider data from 13 hours to 1 hour before the event time. This process is illustrated in Figure 2.

We identified 901 elderly patients from 27,542 adult hospital admissions that spent at least one day in the hospital

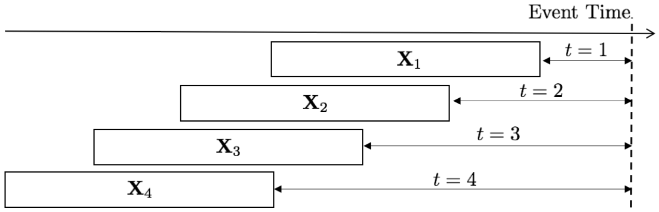

Fig. 2. The data generation process for our study. Each time point only considers data from the 12 hours preceding it.

and had at least one measurement of height. 763 of the 901 patients had at least one observation (out of the 6 clinical measurements) in all 6 time points of interest. Of these 763 patients, 197 of them experienced a cardiac arrest event ($\sim 25.8\%$ prevalence) with the other 566 classified as normal. The data is standardized for each variable across all the patients and possible time points. Therefore the mean and variance of a variable in one of the time points of interest need not necessarily be 0 and 1, respectively. For the missing values, we imputed the values based on the median from patients of the same gender and similar ages, which has been shown to be fairly effective [12].

B. Evaluation Measure

We evaluated the performance of our temporal smoothing model against logistic regression models that were built using each individual time point. Our study used the Monte Carlo cross-validation technique – 10 random training and test splits were created with a 70% - 30% split. We maintained the same prevalence of cardiac arrest patients across the training and test sets. Each model was fit using on the training data and the area under the receiver operating curve (AUC) is measured on the held out 30% test set.

IV. Results

A. Baseline Models

TABLE I

Baseline AUC scores

	Train		Test	
Hour	Mean	SD	Mean	SD
1	0.6588	0.0155	0.6137	0.0289
2	0.6612	0.0116	0.6138	0.0208
3	0.6483	0.0174	0.6068	0.0320
4	0.6777	0.0112	0.6220	0.0230
5	0.6522	0.0125	0.5941	0.0328
6	0.6467	0.0179	0.6306	0.0397

Our baseline models consist of 6 separate logistic regression models, each is trained on data for each individual time point. Table I summarizes the AUC scores across the 10 random splits for both training and test. The results illustrate the difficulty of cardiac arrest prediction with the highest AUC of 0.63. The statistically differences between training and test also suggest that the models may suffer slightly from overfitting. Moreover, we can see that predicting cardiac arrest at least 5 hours prior to the event is an extremely difficult task compared to the other time points.

Figure 1 illustrates the estimated coefficients across the 6 different time points. Although a single split is shown for visualization purposes, the other 9 train–test splits yield similar characteristics. From the figure, we can see that the coefficient values for pulse pressure index, systolic blood pressure, and temperature undergoes drastic changes (including changing signs). Given the time interval is only one hour, one would expect the coefficients from the neighboring points to be similar to one another. Otherwise, how does one explain that if a patient has the same observed pulse pressure index at time point 4, 3 and 2, the measurement goes from a positive effect (time point 4) to negative effect (time point 3) back to a positive effect (time point 2) on the log odds of a cardiac arrest event within the span of 3 hours? This phenomenon suggests an overfitting of the training data at each time point. Moreover, the high variation in the coefficients across adjacent time points makes it difficult to discern a general trend in the coefficient trajectory.

B. Regularization Parameter Effect

We first examine the effect of the temporal smoothing regularization parameter, λ on prediction accuracy. The parameter space is varied evenly on the log scale between 0.01 and 100. Figure 3 illustrates a boxplot of the resulting AUC scores across the 6 time points and 5 regularization parameters. From the figure, we can see that higher regularization generally helps improve the predictive accuracy of the model. This phenomenon is consistent at most of the time points except for 6 hours before the event. At the six hour time point, we see a decrease in the AUC scores, which suggests too much regularization may provide too much smoothing and underfit the data.

In addition to the boxplot of the AUC scores for TTL, the black line in Figure 3 represents the mean test AUC from our baseline model (Table I). The results suggest that TTL-Reg helps improve the predictive performance of our model over the baseline models. The difference in AUC is especially prominent in time points 5 and 6, with a noticeable rise for $\lambda \geq 1$ for cardiac arrest prediction 5 hours before the event. The figure illustrates the potential of TTL-Reg to help build a robust early cardiac arrest prediction model by learning the coefficients from multiple time perspectives.

C. Coefficient Interpretation

Given the results from Figure 3, we chose $\lambda = 1$ for the optimal regularization parameter as it seems to perform comparably across all six different time points. Figure 4 displays the estimated coefficients across the 6 different time points. Although a single split is shown for visualization purposes, the other 9 train–test splits yield similar plots and have been omitted. In contrast to the estimated coefficients for individual learned models at each time point (Figure 1), there is a smoothness to the coefficient trajectories over the different time points. The gradual increase and decrease in the estimated coefficient pathways allow us to discern trends

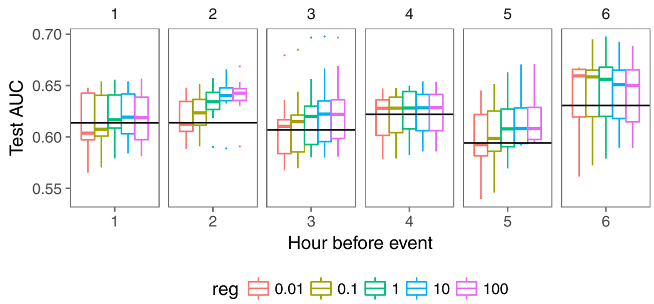

Fig. 3. Boxplot of the AUC scores across the 10 random splits with different regularization values. The black line represents the mean AUC from the baseline model shown in Table I.

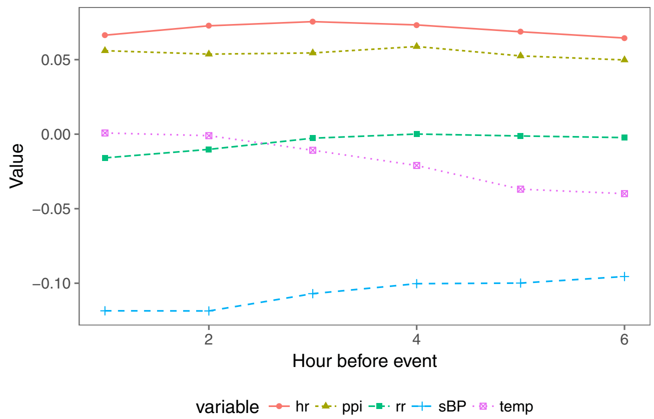

Fig. 4. Trajectories of estimated coefficients for different time points via TTL-reg with $\lambda = 1$.

from the estimated coefficients. For example, we can see that the impact of systolic blood pressure increases the closer we get to the event time. Meanwhile, the respiratory rate which had little effect at time point 6 transitions to a negative effect on the log odds of a cardiac arrest event, while the opposite holds true for a patient's temperature. The figure also shows that the effect of heart rate and pulse pressure index stay fairly constant throughout all the time windows. TTL-Reg not only maintains the model interpretability benefits of logistic regression, but it also can capture temporal trends across different time points which can potentially uncover new knowledge for doctors.

V. DISCUSSION

In this paper, we introduced TTL-Reg, a temporal transfer learning based model to learn a robust cardiac arrest prediction model. The algorithm learns from different time perspectives by smoothing the estimated coefficients of logistic regression from adjacent time points. We show that the parameters of TTL-Reg can be solved iteratively using existing software packages by transforming the objective function into an ℓ_2-regularized logistic regression model. Our model not only yields a coefficient trajectory that can be easily interpreted and potentially uncover new trends but also results in improved early prediction of cardiac arrest patients. The results on 763 ICU patients illustrates the potential of our temporal transfer learning approach.

TTL-Reg can be enhanced from several aspects. From the application perspective, we should examine the coefficient trajectories via a panel of doctors and domain experts to see if it makes medical sense. Our model can also be generalized to more diseases such as predicting other high–mortality events such as septic shock. In addition, it can also be utilized beyond the ICU setting to more generic settings including hospital readmissions and emergency room visits. From the technical aspect, we can explore different forms of regularization including ℓ_1 and ℓ_0 norms. We can also expand the base algorithm to include more complex, non–linear models such as decision trees, support vector machines, and neural networks. Moreover, for these more complicated models, understanding the effect of the model on the interpretability of parameter smoothing is a research topic within itself. Finally, theoretical justification of the temporal smoothing can be explored both from a generative or discriminative model perspective.

REFERENCES

[1] G. I. Fishman, S. S. Chugh, J. P. DiMarco, C. M. Albert, M. E. Anderson, R. O. Bonow, A. E. Buxton, P. S. Chen, M. Estes, X. Jouven, R. Kwong, D. A. Lathrop, A. M. Mascette, J. M. Nerbonne, B. O'Rourke, R. L. Page, D. M. Roden, D. S. Rosenbaum, N. Sotoodehnia, N. A. Trayanova, and Z. J. Zheng, "Sudden Cardiac Death Prediction and Prevention: Report From a National Heart, Lung, and Blood Institute and Heart Rhythm Society Workshop," *Circulation*, vol. 122, no. 22, pp. 2335–2348, Nov. 2010.

[2] C. Sandroni, J. Nolan, F. Cavallaro, and M. Antonelli, "In-hospital cardiac arrest: incidence, prognosis and possible measures to improve survival." *Intensive Care Medicine*, vol. 33, no. 2, pp. 237–245, Feb. 2007.

[3] 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*, vol. 54, no. 2, pp. 125–131, Aug. 2002.

[4] 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*, vol. 65, no. 1, pp. 41–44, Apr. 2005.

[5] 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*, vol. 141, no. 5, pp. 1170–1176, May 2012.

[6] 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*, vol. 8, p. 40, 2011.

[7] J. C. Ho, Y. Park, C. M. Carvalho, and J. Ghosh, "DYNACARE: Dynamic Cardiac Arrest Risk Estimation," in *Proceedings of the Sixteenth International Conference on Artificial Intelligence and Statistics*, 2013, pp. 333–341.

[8] Y. Park, J. C. Ho, and J. Ghosh, "Multivariate temporal symptomatic characterization of cardiac arrest." in *Proceedings of the 35th International Conference of the IEEE Engineering in Medicine and Biology Society*. IEEE, Jan. 2013, pp. 3222–3225.

[9] W. Y. Kim, Y. J. Shin, J. M. Lee, J. W. Huh, Y. Koh, C.-M. Lim, and S. B. Hong, "Modified Early Warning Score Changes Prior to Cardiac Arrest in General Wards," *PLoS ONE*, vol. 10, no. 6, 2015.

[10] S. J. Pan and Q. Yang, "A survey on transfer learning," *IEEE Transactions on Knowledge and Data Engineering*, vol. 22, no. 10, pp. 1345–1359, 2010.

[11] 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*, vol. 39, no. 5, pp. 952–960, May 2011.

[12] J. C. Ho, C. H. Lee, and J. Ghosh, "Septic Shock Prediction for Patients with Missing Data." *ACM Transactions on Management Information System*, vol. 5, no. 1, pp. 1–15, 2014.