


Journal Pre-proof

Contrastive Learning Improves Critical Event Prediction in COVID-19 Patients

Tingyi Wanyan, Hossein Honarvar, Suraj K. Jaladanki, Chengxi Zang, Nidhi Naik, Sulaiman Somani, Jessica K. De Freitas, Ishan Paranjpe, Akhil Vaid, Jing Zhang, Riccardo Miotto, Zhangyang Wang, Girish N. Nadkarni, Marinka Zitnik, Ariful Azad, Fei Wang, Ying Ding, Benjamin S. Glicksberg



PII: S2666-3899(21)00256-7

DOI: <https://doi.org/10.1016/j.patter.2021.100389>

Reference: PATTERN 100389

To appear in: *Patterns*

Received Date: 11 February 2021

Revised Date: 12 September 2021

Accepted Date: 21 October 2021

Please cite this article as: Wanyan T, Honarvar H, Jaladanki SK, Zang C, Naik N, Somani S, De Freitas JK, Paranjpe I, Vaid A, Zhang J, Miotto R, Wang Z, Nadkarni GN, Zitnik M, Azad A, Wang F, Ding Y, Glicksberg BS, Contrastive Learning Improves Critical Event Prediction in COVID-19 Patients, *Patterns* (2021), doi: <https://doi.org/10.1016/j.patter.2021.100389>.

This is a PDF file of an article that has undergone enhancements after acceptance, such as the addition of a cover page and metadata, and formatting for readability, but it is not yet the definitive version of record. This version will undergo additional copyediting, typesetting and review before it is published in its final form, but we are providing this version to give early visibility of the article. Please note that, during the production process, errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

© 2021

Contrastive Learning Improves Critical Event Prediction in COVID-19 Patients

Tingyi Wanyan^{1,2}, Hossein Honarvar¹, Suraj K. Jaladanki¹, Chengxi Zang³, Nidhi Naik¹, Sulaiman Somani¹, Jessica K. De Freitas^{1,4}, Ishan Paranjpe¹, Akhil Vaid¹, Jing Zhang⁵, Riccardo Miotto¹, Zhangyang Wang⁶, Girish N. Nadkarni^{1,7,8}, Marinka Zitnik⁹, Ariful Azad², Fei Wang³, Ying Ding^{10,11}, Benjamin S. Glicksberg^{1,4,12*}

1. Hasso Plattner Institute for Digital Health at Mount Sinai, Icahn School of Medicine at Mount Sinai, New York, NY, USA
2. School of Informatics, Computing, and Engineering, Indiana University, Bloomington, IN, USA
3. Department of Population Health Sciences, Weill Cornell Medicine, New York, NY, USA
4. Department of Genetics and Genomic Sciences, Icahn School of Medicine at Mount Sinai, New York, NY, USA
5. Renmin University of China, Beijing, China
6. Department of Electrical and Computer Engineering, University of Texas at Austin, Austin, TX, USA
7. Division of Nephrology, Department of Medicine, Icahn School of Medicine at Mount Sinai, New York, NY, USA
8. The Charles Bronfman Institute for Personalized Medicine, Icahn School of Medicine at Mount Sinai, New York, NY
9. Department of Biomedical Informatics, Harvard University, USA
10. Dell Medical School, University of Texas at Austin, Austin, TX, USA
11. School of Informatics, University of Texas at Austin, Austin, TX, USA
12. Lead contact

***Correspondence & Lead Contact:** Benjamin S. Glicksberg (benjamin.glicksberg@mssm.edu)

Abstract

Deep Learning (DL) models typically require large-scale, balanced training data to be robust, generalizable, and effective in the context of healthcare. This has been a major issue for developing DL models for the coronavirus-disease 2019 (COVID-19) pandemic where data are highly class imbalanced. Conventional approaches in DL use cross-entropy loss (CEL) which often suffers from poor margin classification. We show that contrastive loss (CL) improves the performance of CEL especially in imbalanced electronic health records (EHR) data for COVID-19 analyses. We use a diverse EHR data set to predict three outcomes: mortality, intubation, and intensive care unit (ICU) transfer in hospitalized COVID-19 patients over multiple time windows. To compare the performance of CEL and CL, models are tested on the full data set and a restricted data set. CL models consistently outperform CEL models with differences ranging from 0.04 to 0.15 for AUPRC and 0.05 to 0.1 for AUROC.

INTRODUCTION

As of May 2021, coronavirus-disease 2019 (COVID-19) has resulted in over 3.4 million reported deaths with over 580,000 occurring in the United States and over 53,000 in New York State¹. Hospital resources such as medication supply and hospital beds quickly became constrained^{2, 3}. Due to the novelty of the COVID-19 pandemic, there is a dearth of relevant data available for research purposes, so electronic health records (EHR) became a valuable tool to study and characterize COVID-19 patients. EHR have already been used in biomedical research for disease and patient clustering, patient trajectory modeling, disease prediction, and clinical decision support, among others⁴. Recent studies discovered important findings to better understand COVID-19 through EHR-based analyses⁵⁻¹¹.

Machine learning (ML) is useful to examine resource allocation and risk stratification of patients, and ML models were successfully used to identify high-risk COVID-19 patients¹²⁻²¹. For temporal data modeling and prediction of patient outcomes in particular, deep learning (DL)²² holds promise over traditional ML techniques that require manual engineering of informative patient features²³⁻²⁵. Heterogeneous graph networks are a powerful DL based graph representation technique that have successfully been utilized in EHR-related studies^{26, 27}. These models have a graphical structure that captures the underlying relationships between disparate medical concepts such as diagnoses and lab tests²⁸. Further, these graph convolutional models can be endowed with an attention mechanism²⁹ to automatically identify how important local neighbors in the graph are to a given medical concept³⁰. Attention models such as RETAIN provide a detailed interpretation of results and maintain high prediction accuracy by using reverse-time attention mechanisms to consolidate past visits³¹.

There are, however, substantial concerns about the limited generalizability of these models in COVID-19 (and in general) as they often underperform in external validation^{32, 33}. Poor generalization of the models is normally due to underspecification³⁴. The underlying aspects of EHR data may also limit model effectiveness³⁵. Healthcare data sets often have inadequate sample sizes in terms of both small hospitals and rare disease populations and exhibit high class imbalance for key outcomes of interest, such as in rare events, as in the case with COVID-19. Several strategies have been utilized to mitigate these data and modeling challenges including up- and down- sampling, pre-training, transfer learning, and federated learning, but each has its limitations for use in EHR research³⁶. Other than these methods, the role of loss function has yet to be thoroughly investigated in the context of COVID-19 EHR work, which is the focus of this work. DL models often use cross-entropy loss (CEL) function, but CEL has been shown to potentially have poor classification margins, making the model less generalizable³⁷. Recently, supervised contrastive loss (CL) has been proposed to improve the classification performance of CEL³⁷. This original CL algorithm used sets of one anchor, many positives, and many negatives to maximize the similarities within the same class. In³⁷, it was shown that CL is more general than triplet and N-pairs losses because for any anchor, all positive pairs including the augmentation of the anchor are used for the loss. In addition, CL has a temperature parameter (τ) in loss that has shown to improve learning. Triplet and N-pairs losses are special cases of supervised CL loss. When one positive and one negative pair are used, CL simplifies to triplet loss^{38, 39}. When positive cases differ from anchor (i.e., excluding augmentations), more than one negative, and no τ are used, CL becomes equivalent to N-pairs loss^{40, 41}.

Though CL has already been applied to learn visual representations of medical images from paired images and text^{42,43}, it has not yet been used for EHR data and COVID-19 in particular. Here, we modify a CL-based architecture from its original formulation for EHR tasks. In the original work, CL was used for representation learning and CEL was used for the downstream classification. While we still utilize CL for patient representation learning, we explicitly incorporate classification objective into our CL loss function. This additional term behaves similar to CEL and maximizes the similarities between patient representation (obtained from sequential models) and outcome representation (obtained from heterogeneous relational models)⁴⁴.

In this study, we compare different sequential models (RETAIN and RNN) models utilizing CEL and our developed CL in predicting critical outcomes of COVID-19 mortality, ICU transfers, and intubation for patients admitted to a large and diverse patient population from five hospitals within the Mount Sinai Health System in New York. Models are tested on a full data set and a restricted data set with severe class imbalance to better elucidate the impacts of these loss functions on model performance. Results are evaluated within the framework of three dimensions: predictive power, patient clustering, and feature importance.

Our main contributions can be summarized as:

1. We propose a novel deep learning framework for EHR data by adding a contrastive regularizer to CEL to improve the performance of prediction tasks.
2. We quantitatively examine the performance of CL for predicting critical events for COVID-19 patients.
3. We show the superior performance of CL framework compared to CEL and traditional ML algorithms, especially when data outcomes are more imbalanced.
4. We provide interpretability of these models that shows how feature importance become more clinically relevant using CL on data sets with imbalanced outcomes.

MATERIALS AND METHODS

Clinical Data and Cohort

We obtained Electrical Health Records (EHR) of COVID-19 patients from five hospitals within the Mount Sinai Healthcare System (MSHS). The collected EHR data contain the following information: COVID-19 status, demographics (age, gender, and race), 55 relevant laboratory test results (listed in Supplementary Table 2), vital signs, specifically heart rate, respiration rate, pulse oximetry, blood pressure (diastolic and systolic), temperature, height, weight, and 12 comorbidities (atrial fibrillation, asthma, coronary artery disease, cancer, chronic kidney disease, chronic obstructive pulmonary disease, diabetes mellitus, heart failure, hypertension, stroke, alcoholism, and liver disease). In addition, we collected information on clinically-relevant outcomes of COVID-19: mortality, discharge, ICU transfer, and intubation. Lab tests and vital signs were measured at multiple time points along the hospital course. In our models, demographics features are used as static features.

Data Pre-processing

We pre-process the vital signs, lab test, and static features by considering the values between 0.5 and 99.5 percentile to remove any inaccurate measurement. For any numerical data, we normalize the data by calculating the standard score (z-score). For categorical data, we use one-hot encoding representation. Numerical data with missing values are included with zeros⁴⁵.

The initial feature input for vital signs at every time step is the vector $X_v \in R^8$ representing the eight features (Figure 1A). Blood pressure is represented as two features: *systolic* and *diastolic*. Similar to vital signs, the initial feature input for lab tests is the vector $X_l \in R^{55}$ representing the 55 features at each time step. If there are more than one vital sign or lab test for each patient, we average the values of the corresponding time step. We concatenate the vital sign vector and lab test vector to form the vector $X_i \in R^{63}$ as the input at every time step -- the subscripts v and l are dropped for simplicity.

Static features consist of demographics (age, gender, and race) and disease comorbidity information, which are detailed in Supplementary Table 2. For Age, we record the normalized numerical value. We represent the Gender feature as a 2 dimensional (male and female) one-hot vector. We use a 5 dimensional one-hot vector to represent the different groups for Race (African American, White, Asian, Other and Unknown). We represent disease comorbidity as a 12 dimensional one-hot vector. We concatenate all the demographic and disease comorbidity features into the vector $X_d \in R^{20}$.

Time Sequence Modeling

To model EHR data as a time sequence, we use a previously developed interpretable predictive model named RETAIN³¹. This model is designed specifically to add feature explainability in terms of feature importance score on time sequence data.

For the RETAIN model, we present each patient for n time steps as in Choi et al.³¹:

$$C_{p,seq} = \sum_{i=1}^n \alpha_i \beta_i \cdot v_i \cdot \# \quad (1)$$

where i denotes the time step, n is total number of time steps, $\alpha_i \in R^1$ is the attention vector of weights, $\beta_i \in R^{l \times 1}$ is the attention weight for each feature, $v_i = W_p X_i$ is a linear projection of the m dimensional input feature, $X_i \in R^m$ using projection matrix $W_p \in R^{l \times m}$.

As another baseline architecture, we utilize a RNN model's $C_{p,seq}$ with the original input feature vector X_i :

$$C_{p,seq} = RNN(X_i) \cdot \# \quad (2)$$

We concatenate the static features vector $C_{p,static}$ with the output of sequential models $C_{p,seq}$ as illustrated in Figure 1A to get the final patient embedding $C_p = C_{p,seq} + C_{p,static}$.

Our RNN architecture is based on long-short term memory (LSTM).

Heterogeneous Relation Modelling

Instead of treating the outcomes as labels, we model them as a directed heterogeneous bipartite graph as illustrated in Figure 1B. We create a triple relationship between patient, outcome, and event, where the outcome is the relation (or edge) between patient node p and event node e . Patient nodes with the same outcome relations are connected to the same event node. Since we are predicting binary outcomes, we have two event nodes representing positive and negative labels. Modeling the data as a bipartite graph provides both label information as well as event and clinical characteristic similarities.

Since the patient and event are two different node types, we use the heterogeneous relational model TransE⁴⁴ to project patient and event node types and their outcome relationships into a shared latent space. The TransE model aims to relate different type of nodes by their relationship type and represents the relationship type (outcome) as a translation vector between the two node types. This relation is expressed as:

$$\begin{aligned} \hat{C}_e &= \delta(X_e W_e + b_e) \\ C_e &= \hat{C}_e - R_o, \# (3) \end{aligned}$$

where X_e is the binary outcome representation, $W_e \in R^{2 \times l}$ are learnable projection parameters for latent dimension l , $b_e \in R^l$ are bias parameters, and δ is a non-linear activation function. We use a two dimensional vector $X_e \in (0,1)^2$ to represent positive or negative outcomes. \hat{C}_e is the latent representation of outcomes, C_e is translated representation from \hat{C}_e in the projection space by the learnable translation relational vector R_o , which is the relation vectors representing outcome relation that connects patients to positive event and negative event nodes, respectively.

After the projection, we apply similarity comparisons between these two representations (C_p and C_e) in the shared latent space.

Loss Function Delineation

After the assembly of the bipartite relational graph, we aim to predict the binary outcome of a patient by maximizing the similarity between the binary outcome latent representation and patient representation. The bipartite relational graph also considers the similarities within patient latent representations that connect to the same outcome. Therefore, the objective function is expressed as:

$$L = P(N_c(u)|C_p(u)) , \# (4)$$

where u is the patient node of interest in training sample, $C_p(u)$ is the latent representation of patient node u . $N_c(u)$ is the patient node's neighboring nodes which consist of binary outcome node representation C_e and similar patient nodes representations C_p . In Eq. (4), we optimize the proximity between the representations of center patient node and its neighboring nodes.

The similarity between these latent representations is represented as an inner-product, and we directly apply noise contrastive estimation (NCE) loss to capture the condition probability in Eq. (4)⁴⁶.

$$\mathcal{L} = - \sum_{u \in V} \left[\sum_{c_i \in N_c(u)} \log \sigma(\vec{c}_i \cdot \vec{u}) + \sum_{j=1}^K E_{c_j \sim P_v(c_j)} \log \sigma(-\vec{c}_j \cdot \vec{u}) \right], \# (5)$$

where V is the training patients data set, \vec{c}_i is the latent representation vector of j -th context node in $N_c(u)$, \vec{u} is the center node latent representation. K is the number of negative samples, $P_v(c_j)$ is the negative sampling distribution. $\vec{c}_j \cdot \vec{u}$ are the co-occurrence positive representation pairs, and $-\vec{c}_j \cdot \vec{u}$ are the negative sampling pairs. The non-linear function $\sigma(x) = \frac{1}{1 + \exp(-x)}$ captures the similarity score between the representation pairs.

We rewrite the objective function in Eq. (5) in our own notations as follows:

$$\begin{aligned} \mathcal{L} &= - \sum_{u \in V} \left[\sum_{c_p(j) \in N_c(u)} \log \sigma(C_e(u) \cdot C_p(u)) + \log \sigma(-C_e^*(u) \cdot C_p(u)) + \right. \\ &\quad \left. \sum_{c_p(j) \in N_c(u)} \log \sigma(C_p(j) \cdot C_p(u)) + \sum_{j=1}^K E_{C_p^*(j) \sim P_v(C_p^*(j))} \log \sigma(-C_p^*(j) \cdot C_p(u)) \right], \#(6) \end{aligned}$$

where C_e is the projected latent representation of the binary outcome node that connects to a given patient representation C_p , and the inner product between C_e and C_p measures the similarity between these two representations. superscripts * shows the opposite outcome node of C_e and does not connect to the patient of interest u . $C_p(j)$ is the similar context patient node representations that connects to the same outcome node as the patient node u . $C_p^*(j)$ are the context patient node representations that connects to the opposite outcome node as the patient node.

The first two terms in Eq. (6) capture the label information between outcome with the patient node, so they function as cross-entropy loss. In the Supplementary Materials, we prove the lemma that the first two terms in Eq. (6) are equivalent to the cross-entropy loss in general and therefore the improvements are primarily due to CL inter-pateint terms. The last two terms provide additional information of similar patients that connect to the same outcome, and dissimilar patient information that connect to the other outcome.

We set a weight factor α to weigh the importance of the last two parts of Eq.(6) that captures similar patient information, which is the main improvement due to CL. Our final objective function is as follows:

$$L = L_{ep} + L_{ep}^* + \alpha(L_{pp} + L_{pp}^*)$$

$$\begin{aligned}
\mathcal{L}_{ep} &= - \sum_{u \in V} \log \sigma \left(C_e(u) \cdot C_p(u) \right) \\
\mathcal{L}_{ep}^* &= - \sum_{u \in V} \log \sigma \left(-C_e^*(u) \cdot C_p(u) \right) \\
\mathcal{L}_{pp} &= - \sum_{u \in V} \sum_{C_p(j) \in N_c(u)} \log \sigma \left(C_p(j) \cdot C_p(u) \right) \\
\mathcal{L}_{pp}^* &= - \sum_{u \in V} \sum_{j=1}^K E_{C_p^*(j) \sim P_v(C_p^*(j))} \log \sigma \left(-C_p^*(j) \cdot C_p(u) \right) . \# (7)
\end{aligned}$$

In this work, we use the optimal $\alpha = 0.8$ that achieved the best performance for all models and tasks. We describe the experiments for exploring different values of α in Supplementary Table 1.

After minimizing the CL from Eq. (7), we obtain learned latent representation for events C_e , which are used to predict the probability of the events as follows:

$$P \left(C_p(u) \right) = \sigma \left(C_e(u) \cdot C_p(u) \right) , \# (8)$$

where Y_e represents the logit prediction for positive outcomes (mortality, intubation, and ICU transfer).

Feature Importance Scoring

The linear projection matrix W_p from the RETAIN model allows to interpret variable importance at each time step. Our goal is to predict the probability of the outcome given a center patient representation. We can write this probability the same as Eq. (8).

We can combine Eq. (1), Eq. (3) and Eq. (8) to derive the similarity score as follows:

$$\begin{aligned}
P \left(Y_e \mid C_p(u) \right) &\propto (W_e X_e + b_e - R_o)^T \left(\sum_{i=1}^n \alpha_i \beta_i \odot W_p X_i \right) = \\
&\sum_{i=1}^n \left(\alpha_i \beta_i \odot (X_e^T W_e^T + b_e^T - R_o^T) W_p \right) X_i .
\end{aligned}$$

The contribution score for a specific feature k at time step i for input sample is derived as:

$$\omega(Y_e, X_{i,k}) = (\alpha_i \beta_i \odot (X_e^T W_e^T + b_e^T - R_o^T) W_p) X_{i,k} . \# (9)$$

This is the similarity score between the positive outcome latent representation and patient latent representation for a RETAIN model with CL loss. The larger values of ω indicates that the feature k has a large contribution towards the prediction result.

For interpretability of a RETAIN model with CEL, we directly compute the importance score in a similar manner as in Choi et al.³¹:

$$\omega(y_m, x_{i,k}) = \alpha_i W_c (\beta_i \odot W_p) X_{i,k} , \# (10)$$

where y_m is the label for m -th sample.

Baselines

In this work, we are evaluating the performance of CL using two time sequence models (RETAIN+CL and RNN+CL). As baseline models, we use the CEL with the same time sequence models (RETAIN+CE and RNN+CE) to evaluate the potential improvements of CL.

As reference and comparison with the CL, the objective function for CEL is as follows:

$$L_{CEL} = -\frac{1}{N} \sum_{m=1}^N (y_m \log \log (\hat{y}_m) + (1 - y_m) \log \log (1 - \hat{y}_m)) , \# (11)$$

where the logit output for the m -th sample is:

$$\hat{y}_m = \text{Sigmoid}(W_c C_p + B_c) , \# (12) \text{ where } C_p \text{ is the latent patient representation, } W_e \in R^{2 \times l} \text{ and } B_c \in R^2 \text{ are the binary projection and bias parameters.}$$

We also compare the performance of the sequential models with respect to four traditional ML algorithms: logistic regression (LG), random forest (RF), support vector machine (SVM); XGBoost (XGB)⁴⁷.

Experiment Design

We perform three prediction tasks: mortality, ICU transfer, and intubation. We train our models on predicting events for two-time frames: 24 and 48 hours before the occurrence of a binary outcome. Longitudinal data (lab tests and vital signs) are binned within windows of 6 hours and averaged if there are more than one measurement per window. The time binning representation is illustrated in Figure 1C. In the training phase, for each positive outcome, we utilize the exact time of an event to generate time frames for the experiments. For patients that did not have these

outcome events, we needed a representative frame of reference to align against. Therefore, we compute the mean and standard deviation for the length of time that elapsed from admission for all the affiliated outcomes independently. For patients without an event, we randomly pick a time to use as a reference end point using a Gaussian distribution with the mean and standard deviation obtained from the positive training data. This procedure is shown in Figure 1D.

We also perform the comparisons of outcomes for both full sample and in artificial scenarios of more extreme imbalance (i.e., restricted sample) to determine the extent of performance differences between the two loss functions. In terms of generating the samples for any prediction using CEL, we find that a minimum of 5% positive labels is required to detect both negatives and positives. Therefore, we choose any positive labels percentage to be greater than 5%. For the experiments with full sample, we use all available data where the percentages of positive labels is 23% mortality, 17% ICU transfer, and 10% intubation. For the experiments with restricted sample, we perform down-sampling to reduce the percentage of positive labels by randomly removing a percentage of positive labels. The percentages of positive labels in the down sampled data set are 7% mortality, 7% ICU transfer, and 5% intubation. In the restricted sample, the percentage of each positive label is less than half of the original positive label percentage.

Training details

In the training stage, after picking a patient node u , we select the binary outcome node that connects to u . Then we uniformly pick m similar patient nodes that also connect to this outcome node, and we use these samples as positive training pairs. For negative sampling, we first pick the binary outcome node that does not connect to the patient node u , and we then uniformly pick q similar patient nodes that connect to this outcome node. We utilize these samples as the negative training pairs. In this work, we use $m = 2$ and $q = 1$ to prioritize the positive samples.

For validation purposes, we perform 10-fold cross validation for the patient of interest and record the mean evaluation values across 10 folds to determine the performance of CL model against the CE model.

All models and data sets are evaluated using the following metrics: area under the receiver-operating characteristic (AUROC), area under the precision, and recall curve (AUPRC). It is important to note that AUPRC is a more reliable metric for imbalanced samples because that takes into account NPV and PPV⁴⁸.

Results

Model Comparison and Evaluation of Predictive Performance

We first evaluate the results between loss functions using all data (i.e., full sample) for all three tasks, namely mortality, ICU transfer, and intubation prediction. The ROC and PR curves for these cross validated results are shown in Fig.2,3(A-C) and metrics are tabulated in Table 1. For mortality prediction (23% positive label percentage), we observe that the AUROC and AUPRC scores are similar between CEL and CL. For RETAIN model, under 24h prediction time frame, the AUROC and AUPRC scores are $0.92 \pm 0.01, 0.82 \pm 0.02$ for CEL and $0.92 \pm 0.01, 0.84 \pm$

0.02 for CL. These results indicate that two loss models achieve similar performance when we have relatively more balanced label ratios. For intubation prediction, the positive label percentage is 10.7 %, which is less than half of the mortality labels. We observe relatively larger performance increases of 0.03 ± 0.02 for AUROC score and 0.09 ± 0.02 for AUPRC in CL compared to CEL for the RETAIN model. For ICU transfer prediction (17% positive label percentage) using RETAIN, we observe that the AUROC, AUPRC performance of CL is around 0.84 ± 0.01 , 0.60 ± 0.02 which are slightly higher than CEL performance of 0.81 ± 0.01 , 0.57 ± 0.02 .

To further evaluate the above trends, we conducted an additional experiment to assess how CEL and CL functions perform on the same tasks with more imbalanced outcome ratios. This scenario may be the case in smaller hospital cohorts and for other outcomes. We perform random down-sampling on positive labels (i.e., restricted sample): for mortality, intubation, and ICU transfer prediction tasks, we randomly down-sampled the positive labels to 7%, 5%, and 7%, respectively. The ROC and PR curves are shown in Fig.2,3(E-F) and the performance metrics are recorded in Table 2. The results of the experiment consistently show lower performance using CEL compared to CL. For RETAIN model under 24h time frame, AUROC and AUPRC values are higher for CL than CEL for all outcomes. For instance, the AUROC increases from 0.80 ± 0.02 (CEL) to 0.88 ± 0.02 (CL) and AUPRC increases from 0.35 ± 0.03 (CEL) to 0.45 ± 0.02 (CL) for intubation prediction. These findings show that under cases with extremely unbalanced label data, models using CL perform better than using CEL.

Our results show that CEL and CL have competitive performance in the full data set but we generally find that for all tasks in restricted data set, models with CEL have lower performance compared to CL (Table 2). Other than the 24-hour time frame, we also perform the same exact analysis for the 48-hour time frame, which shows consistent trends. Finally, we use the RNN as the baseline model other than RETAIN model for all the predictions described earlier and our conclusions hold true (see Table 2).

We have also provided the results of using traditional ML algorithms in both tables. Our results generally show the superiority of DL with respect to these baselines.

Identification of Important Clinical Features

We assess the ability of CEL and CL functions in our models to identify relevant features of interest in full and restricted data sets. Specifically, we performed feature importance score calculations for the RETAIN model on predicting mortality as a representative case.

First, we generated the feature importance scores comparing two models (RETAIN-CL against RETAIN-CEL) over the 24-hour time frame for the full data set for mortality prediction. Fig. 5A and B shows the normalized feature importance heat maps for CL and CEL with the columns representing four 6-hour windows. These heat maps display similar importance scores in terms of key features and their magnitudes. RETAIN identified lab tests and vital sign features that are considered important by both loss models: pulse oximetry (0.26 for CEL, 0.21 for CL); aspartate aminotransferase (0.69 for CEL, 0.72 for CL); blood urea nitrogen (0.12 for CEL, 0.12 for CL); lactate (0.24 for CEL, 0.23 for CL); lactate dehydrogenase (0.28 for CEL, 0.3 for CL). All these parameters indicate important aspects of an ill COVID-19 patient.

We then generate feature importance scores using two loss functions for mortality prediction in the restricted data set (i.e., down-sampling the positive label to 7%) for the RETAIN model. We assess how the variable importance score changes under these different conditions for different loss functions. The corresponding heat map is plotted in Fig. 5C and D. We observe that using CL can still capture the highly scored features identified in the full data set (i.e., weigh similar key features). On the other hand, CEL fails to capture some important features. Of particular interest, the importance of pulse oximetry is no longer prioritized in the restricted sample using CL (importance value is 0.09 for CEL comparing to 0.35 for CL). Also, blood urea nitrogen, lactate have lower importance values of 0.02, 0.09 for CEL compared to 0.15, 0.36 for CL. These findings reaffirm our hypothesis with CL are more robust when the outcome labels are highly imbalanced.

Visualizing Patient Embeddings

Finally, we generated 2D t-SNE projections to compare patient embedding representations for RETAIN models between the CL and CEL in predicting all three medical events within 24-hour intervals and the results are shown in Fig. 4. The first two columns show patient embeddings using the full sample data set, and the last two columns show embeddings for the restricted data set. Blue dots represent positive labels and red dots represent negative labels. Models with both loss functions show clear clustering of positive and negative labels in the full data sets. However, when the data set is restricted (i.e., lowered positive labels), the model with CEL consistently show less clear patterns and poorer clustering of patients for all outcomes. In contrast, the RETAIN model with CL maintains its ability to group patients by their clinical outcomes.

DISCUSSION AND CONCLUSION

In this work, we develop a new DL model based on the CL for predicting three critical events (mortality, intubation, and ICU transfer) for COVID-19 patients using EHR data. Most DL-based EHR analyses utilize CEL as part of modeling. To the best of our knowledge, this is one of the first studies to demonstrate the utility of CL in EHR predictive analyses. We demonstrate the benefit of CL in multiple tasks with imbalanced outcome labels, which is particularly pertinent in the context of COVID-19. We compare the performance of different sequential architectures (RNN and RETAIN) for both CL and the conventional CEL model under two-time window horizons. We also compare the performance of our developed framework with respect to traditional ML baselines. We conduct further experiments for each outcome in a restricted data set of even more imbalanced outcomes and show the benefit of CL is even more pronounced via three separate experimental tasks, namely predictive performance, feature importance, and clustering.

The observed improvements in predictions come from the specific form of CL, which does not only maximize the similarity between patient and outcome embedding representations, but also maximizes the similarities between patient representations related to a specific outcome. However, CEL mainly focuses on maximizing the similarity between patient and outcome one-hot representations. Therefore, CL tends to maximize the margins between classes better than CEL. The better margin classification ability of CL lead to higher performance for imbalanced data, for which the classification is difficult due to poor margin classification of CEL in addition to differences in data distributions for each class.

Our study contains several limitations which need to be addressed. First, we only compare CL to one other loss function, although it is widely used. Second, we only assess two sequential modeling

techniques. Another main limitation of our study is excluding several laboratory values due to high levels of missingness, which may impact some of our interpretations. Furthermore, we utilize a specific time sequence modeling representation. Lastly, our work focuses only on one disease use-case. For future studies, other modeling architecture can be assessed to evaluate the generalizability of our CL approach. Also, additional analysis is required to understand feature importance differences between the loss functions and to determine if our CL methodology is applicable to other healthcare data sets. We also plan to assess performance of this strategy for predictive tasks in other diseases such as acute kidney injury. Another important investigation will be long-term predictions other than 24 and 48 hours windows studied in this work for longer-term critical events. We believe this work represents an effective demonstration of the power of using CL for machine learning work for predictive tasks using EHRs.

Table 1: Binary Outcome Prediction Performance on Different Models using **Full Sample Data** (positive label percentage: 23% for mortality, 10% for intubation, 17% for ICU transfer). CEL: cross-entropy loss; CL: contrastive loss; LG: logistic regression, RF: random forest, SVM: support vector machine; XGB: XGBoost. All predictions are calculated from 10-fold cross validation, for which we record the mean value and standard deviation as confident intervals across folds. Bold values represent best model performance per event.

Time Window	Event	Model Type	AUROC	AUPRC
24h	Mortality	LG	0.85(0.02)	0.65(0.01)
		RF	0.82(0.02)	0.63(0.01)
		SVM	0.79(0.02)	0.61(0.02)
		XGB	0.83(0.02)	0.65(0.02)
		RNN+CE	0.91(0.01)	0.82(0.01)
		RETAIN+CE	0.92(0.01)	0.82(0.02)
		RNN+CL	0.91(0.02)	0.83(0.01)
		RETAIN+CL	0.92(0.01)	0.84(0.02)
	Intubation	LG	0.82(0.01)	0.47(0.02)
		RF	0.81(0.01)	0.46(0.01)
		SVM	0.76(0.02)	0.42(0.01)
		XGB	0.78(0.02)	0.46(0.01)

		RNN+CE	0.83(0.02)	0.49(0.02)	
		RETAIN+CE	0.85(0.02)	0.48(0.02)	
		RNN+CL	0.91(0.02)	0.56(0.02)	
		RETAIN+CL	0.91(0.02)	0.56(0.03)	
	ICU Transfer	LG	0.81(0.01)	0.52(0.01)	
		RF	0.82(0.02)	0.55(0.01)	
		SVM	0.79(0.01)	0.52(0.01)	
		XGB	0.78(0.02)	0.55(0.01)	
		RNN+CE	0.83(0.01)	0.57(0.02)	
		RETAIN+CE	0.81(0.02)	0.57(0.02)	
		RNN+CL	0.86(0.01)	0.62(0.02)	
		RETAIN+CL	0.85(0.01)	0.59(0.02)	
	48h	Mortality	LG	0.85(0.01)	0.64(0.02)
			RF	0.81(0.01)	0.61(0.02)
SVM			0.81(0.01)	0.63(0.01)	
XGB			0.88(0.01)	0.69(0.01)	
RNN+CE			0.90(0.02)	0.82(0.02)	
RETAIN+CE			0.92(0.01)	0.83(0.01)	
RNN+CL			0.92(0.02)	0.82(0.01)	
RETAIN+CL			0.93(0.01)	0.84(0.01)	
Intubation		LG	0.79(0.01)	0.45(0.02)	
		RF	0.77(0.02)	0.44(0.01)	
	SVM	0.73(0.01)	0.39(0.01)		

		XGB	0.82(0.01)	0.49(0.02)
		RNN+CE	0.69(0.04)	0.40(0.03)
		RETAIN+CE	0.78(0.03)	0.39(0.03)
		RNN+CL	0.86(0.03)	0.54(0.02)
		RETAIN+CL	0.93(0.01)	0.51(0.03)
	ICU Transfer	LG	0.79(0.02)	0.50(0.01)
		RF	0.81(0.01)	0.54(0.01)
		SVM	0.77(0.02)	0.49(0.02)
		XGB	0.81(0.01)	0.57(0.02)
		RNN+CE	0.80(0.01)	0.54(0.02)
		RETAIN+CE	0.81(0.02)	0.52(0.02)
		RNN+CL	0.83(0.02)	0.60(0.02)
		RETAIN+CL	0.83(0.01)	0.59(0.02)

Table 2: Binary Outcome Prediction Performance on Different Models using **Restricted Sample Data** (positive label percentage: 7% for mortality, 5% for intubation, 7% for ICU transfer). CE: cross-entropy loss; CL: contrastive loss LG: logistic regression, RF: random forest, SVM: support vector machine; XGB: XGBoost. All predictions are calculated from 10-fold cross validation, for which we record the mean value and standard deviation as confident intervals across folds. Bold values represent best model performance per event.

Time Window	Event	Model Type	AUROC	AUPRC
	Mortality	LG	0.78(0.02)	0.51(0.02)
		RF	0.61(0.02)	0.24(0.03)
		SVM	0.60(0.02)	0.22(0.02)
		XGB	0.65(0.02)	0.25(0.01)
		RNN+CE	0.83(0.03)	0.53(0.05)

24h		RETAIN+CE	0.86(0.02)	0.50(0.05)
		RNN+CL	0.91(0.03)	0.62(0.04)
		RETAIN+CL	0.91(0.01)	0.59(0.05)
	Intubation	LG	0.76(0.02)	0.33(0.02)
		RF	0.75(0.01)	0.34(0.01)
		SVM	0.75(0.02)	0.32(0.02)
		XGB	0.77(0.02)	0.39(0.02)
		RNN+CE	0.79(0.03)	0.35(0.03)
		RETAIN+CE	0.80(0.02)	0.35(0.03)
		RNN+CL	0.88(0.02)	0.48(0.03)
		RETAIN+CL	0.88(0.02)	0.45(0.02)
	ICU Transfer	LG	0.77(0.01)	0.39(0.02)
		RF	0.74(0.01)	0.36(0.02)
		SVM	0.76(0.02)	0.38(0.01)
		XGB	0.78(0.02)	0.36(0.01)
		RNN+CE	0.78(0.01)	0.41(0.02)
		RETAIN+CE	0.76(0.03)	0.43(0.04)
		RNN+CL	0.86(0.02)	0.53(0.03)
		RETAIN+CL	0.85(0.02)	0.51(0.03)
	Mortality	LG	0.77(0.02)	0.33(0.02)
RF		0.57(0.02)	0.12(0.02)	
SVM		0.62(0.02)	0.19(0.03)	
XGB		0.69(0.02)	0.29(0.02)	

48h		RNN+CE	0.85(0.03)	0.55(0.04)
		RETAIN+CE	0.90(0.03)	0.53(0.03)
		RNN+CL	0.92(0.03)	0.63(0.04)
		RETAIN+CL	0.91(0.02)	0.64(0.04)
	Intubation	LG	0.79(0.01)	0.36(0.02)
		RF	0.78(0.02)	0.35(0.01)
		SVM	0.79(0.01)	0.34(0.01)
		XGB	0.82(0.02)	0.40(0.01)
		RNN+CE	0.70(0.03)	0.34(0.04)
		RETAIN+CE	0.74(0.02)	0.31(0.03)
		RNN+CL	0.83(0.02)	0.44(0.02)
		RETAIN+CL	0.85(0.02)	0.44(0.03)
	ICU Transfer	LG	0.79(0.02)	0.41(0.01)
		RF	0.75(0.01)	0.35(0.01)
		SVM	0.77(0.01)	0.37(0.02)
		XGB	0.79(0.01)	0.41(0.02)
		RNN+CE	0.72(0.04)	0.38(0.03)
		RETAIN+CE	0.75(0.04)	0.43(0.04)
		RNN+CL	0.82(0.01)	0.51(0.02)
		RETAIN+CL	0.82(0.02)	0.50(0.03)

EXPERIMENTAL PROCEDURE**Resource Availability***Lead Contact*

Any further information, question or request should be sent to Benjamin S. Glicksberg (benjamin.glicksberg@mssm.edu)

Materials Availability

Our study did not involve any physical materials.

Data and Code Availability

Our data is not available due to institutional review board (IRB) rules for privacy protection. For reproducibility, our code is available at https://github.com/Tingyiwanyan/CL_covid and runs with TensorFlow 1.15.

DECLARATION OF INTEREST

The authors declare no competing interest

Figure Captions

Figure 1: Data and modeling schemas. A) Architecture with Contrastive Loss. Electronic health record (EHR) data are modeled to create patient and event embedding representations which are fed into our contrastive loss (CL) equation. B) Representation Space. CL simultaneously pushes positive patients and event embeddings (i.e., concordant with respect to the outcome of the patient of interest, respectively) away from negative ones. C) Time Binning. Schematic to visualize how we model time sequence. We have two outcome windows (i.e., 24- and 48-hours prior to event) and bin data by 6-hour chunks. D): Selection of Event Timing for Null Outcomes. For patients that do not experience the outcome of interest, we generate a data-driven event time to align against as in C. We compute the mean and standard deviation for the length of time that elapsed from admission for all patients with the affiliated outcomes independently. For patients without an event, we randomly pick a time to use as a reference end point using a Gaussian distribution with the mean and standard deviation obtained from the positive training data.

Figure 2: Receiver Operator Curves for All Predictive Tasks in a 24-hour Time Frame. Performance is assessed for both contrastive loss (CL) and cross-entropy loss (CEL) for both RNN and RETAIN modeling strategies. A) Mortality with full data set (23% positive labels); B) Intubation with full data set (11% positive labels); C) ICU transfer with full data set (17% positive labels); D) Mortality with restricted data set (7% positive labels); E) Intubation with restricted data set (5% positive labels); F) ICU transfer with restricted data set (7% positive labels).

Figure 3: Precision-Recall (PR) Curves for All Event Predictions a 24-hour Time Frame. Performance is assessed for both contrastive loss (CL) and cross-entropy loss (CEL) for both RNN and RETAIN modeling strategies. A) Mortality with full data set (23% positive labels); B) Intubation with full data set (11% positive labels); C): ICU transfer with full data set (17% positive labels); D): Mortality with restricted data set (7% positive labels); E Intubation with restricted data set (5% positive labels); F) ICU transfer with restricted data set (7% positive labels).

Figure 4: t-SNE Latent Embedding Comparisons for All Event Prediction within 24-hour Time Frame Using RETAIN. Blue dots represent positive labels and red dots represent negative labels. The plot is organized by outcome per row namely 1st: mortality, 2nd: intubation, and 3rd: ICU

transfer. The first and third columns represent contrastive loss (CL) plots and the second and forth represent cross-entropy loss (CEL). A) Mortality prediction with CL for the full data set (23% positive labels); B) Mortality prediction with CL for the full data set; C) Mortality prediction with CL for the restricted data set (7% Positive Labels); D) Mortality prediction with CEL for the restricted data set. E) Intubation prediction with CL for the full data set (10% positive labels); F) Intubation prediction with CEL for the full data set; G) Intubation prediction with CL for the restricted data set (5% positive labels); H) Intubation prediction with CEL for the restricted data set. I) ICU transfer prediction with CL for the full data set (17% positive labels); J) ICU transfer prediction with CEL for the full data set; K) ICU transfer prediction with CL for the restricted data set (7% positive labels); L) ICU transfer prediction with CEL for the restricted data set.

Acknowledgments

We thank the Clinical Data Science and Mount Sinai Data Warehouse teams for providing data. We appreciate all of the providers who contributed to the care of these patients. This work was supported by U54 TR001433-05, National Center for Advancing Translational Sciences, National Institutes of Health. YD acknowledges support from Amazon Machine Learning Research Award. FW and CZ acknowledge the support from NSF 1750326 and 2027970, as well as the AWS Machine Learning for Research Award and Google Faculty Research Award.

Author Contributions

TW, FW, AZ, YD, and BSG designed the study. TW, HH, and CZ performed the analyses. TW, SKJ, NN, SS, JKDF, IP, AV, RM, and GNN processed the data. TW, HH, RM, MZ, ZW, AA, FW, YD, and BSG interpreted the results. AA, FW, YD, and BSG supervised the work. All authors edited and approved the final version of the manuscript.

Declaration of Interests

The authors declare no competing interests.

References

- [1] Dong E, Du H, Gardner L: An interactive web-based dashboard to track COVID-19 in real time. *Lancet Infect Dis* 2020, 20:533-4.
- [2] Thompson CN, Baumgartner J, Pichardo C, Toro B, Li L, Arciuolo R, Chan PY, Chen J, Culp G, Davidson A, Deviney K, Dorsinville A, Eddy M, English M, Fireteanu AM, Graf L, Geevarughese A, Greene SK, Guerra K, Huynh M, Hwang C, Iqbal M, Jessup J, Knorr J, Latash J, Lee E, Lee K, Li W, Mathes R, McGibbon E, McIntosh N, Montesano M, Moore MS, Murray K, Ngai S, Paladini M, Paneth-Pollak R, Parton H, Peterson E, Pouchet R, Ramachandran J, Reilly K, Sanderson Slutsker J, Van Wye G, Wahnich A, Winters A, Layton M, Jones L, Reddy V, Fine A: COVID-19 Outbreak - New York City, February 29-June 1, 2020. *MMWR Morb Mortal Wkly Rep* 2020, 69:1725-9.
- [3] McMahon DE, Peters GA, Ivers LC, Freeman EE: Global resource shortages during COVID-19: Bad news for low-income countries. *PLoS Negl Trop Dis* 2020, 14:e0008412.
- [4] Glicksberg BS, Johnson KW, Dudley JT: The next generation of precision medicine: observational studies, electronic health records, biobanks and continuous monitoring. *Hum Mol Genet* 2018, 27:R56-R62.

- [5] Clifford CT, Pour TR, Freeman R, Reich DL, Glicksberg BS, Levin MA, Klang E: Association between COVID-19 diagnosis and presenting chief complaint from New York City triage data. *Am J Emerg Med* 2020.
- [6] Reeves JJ, Hollandsworth HM, Torriani FJ, Taplitz R, Abeles S, Tai-Seale M, Millen M, Clay BJ, Longhurst CA: Rapid response to COVID-19: health informatics support for outbreak management in an academic health system. *J Am Med Inform Assoc* 2020, 27:853-9.
- [7] Somani SS, Richter F, Fuster V, De Freitas JK, Naik N, Sigel K, Mount Sinai CIC, Bottinger EP, Levin MA, Fayad Z, Just AC, Charney AW, Zhao S, Glicksberg BS, Lala A, Nadkarni GN: Characterization of Patients Who Return to Hospital Following Discharge from Hospitalization for COVID-19. *J Gen Intern Med* 2020, 35:2838-44.
- [8] Wagner T, Shweta F, Murugadoss K, Awasthi S, Venkatakrisnan AJ, Bade S, Puranik A, Kang M, Pickering BW, O'Horo JC, Bauer PR, Razonable RR, Vergidis P, Temesgen Z, Rizza S, Mahmood M, Wilson WR, Challener D, Anand P, Liebers M, Doctor Z, Silvert E, Solomon H, Anand A, Barve R, Gores G, Williams AW, Morice WG, 2nd, Halamka J, Badley A, Soundararajan V: Augmented curation of clinical notes from a massive EHR system reveals symptoms of impending COVID-19 diagnosis. *Elife* 2020, 9.
- [9] Wang Z, Zheutlin A, Kao YH, Ayers K, Gross S, Kovatch P, Nirenberg S, Charney A, Nadkarni G, De Freitas JK, O'Reilly P, Just A, Horowitz C, Martin G, Branch A, Glicksberg BS, Charney D, Reich D, Oh WK, Schadt E, Chen R, Li L: Hospitalised COVID-19 patients of the Mount Sinai Health System: a retrospective observational study using the electronic medical records. *BMJ Open* 2020, 10:e040441.
- [10] Williamson EJ, Walker AJ, Bhaskaran K, Bacon S, Bates C, Morton CE, Curtis HJ, Mehrkar A, Evans D, Inglesby P, Cockburn J, McDonald HI, MacKenna B, Tomlinson L, Douglas IJ, Rentsch CT, Mathur R, Wong AYS, Grieve R, Harrison D, Forbes H, Schultze A, Croker R, Parry J, Hester F, Harper S, Perera R, Evans SJW, Smeeth L, Goldacre B: Factors associated with COVID-19-related death using OpenSAFELY. *Nature* 2020, 584:430-6.
- [11] Wu Z, McGoogan JM: Characteristics of and Important Lessons From the Coronavirus Disease 2019 (COVID-19) Outbreak in China: Summary of a Report of 72314 Cases From the Chinese Center for Disease Control and Prevention. *JAMA* 2020, 323:1239-42.
- [12] Yang HS, Hou Y, Vasovic LV, Steel PA, Chadburn A, Racine-Brzostek SE, Velu P, Cushing MM, Loda M, Kaushal R: Routine laboratory blood tests predict SARS-CoV-2 infection using machine learning. *Clinical chemistry* 2020, 66:1396-404.
- [13] Wynants L, Van Calster B, Collins GS, Riley RD, Heinze G, Schuit E, Bonten MM, Dahly DL, Damen JA, Debray TP: Prediction models for diagnosis and prognosis of covid-19: systematic review and critical appraisal. *bmj* 2020, 369.
- [14] Liang W, Liang H, Ou L, Chen B, Chen A, Li C, Li Y, Guan W, Sang L, Lu J: Development and validation of a clinical risk score to predict the occurrence of critical illness in hospitalized patients with COVID-19. *JAMA Internal Medicine* 2020.
- [15] Gysi DM, Valle ÍD, Zitnik M, Ameli A, Gan X, Varol O, Sanchez H, Baron RM, Ghiassian D, Loscalzo J: Network medicine framework for identifying drug repurposing opportunities for covid-19. *arXiv preprint arXiv:200407229* 2020.
- [16] Yan L, Zhang H-T, Goncalves J, Xiao Y, Wang M, Guo Y, Sun C, Tang X, Jing L, Zhang M: An interpretable mortality prediction model for COVID-19 patients. *Nature Machine Intelligence* 2020:1-6.
- [17] Vaishya R, Javaid M, Khan IH, Haleem A: Artificial Intelligence (AI) applications for COVID-19 pandemic. *Diabetes & Metabolic Syndrome: Clinical Research & Reviews* 2020.

- [18] Mei X, Lee H-C, Diao K-y, Huang M, Lin B, Liu C, Xie Z, Ma Y, Robson PM, Chung M: Artificial intelligence-enabled rapid diagnosis of patients with COVID-19. *Nature Medicine* 2020:1-5.
- [19] Vaid A, Somani S, Russak AJ, De Freitas JK, Chaudhry FF, Paranjpe I, Johnson KW, Lee SJ, Miotto R, Richter F: Machine Learning to Predict Mortality and Critical Events in a Cohort of Patients With COVID-19 in New York City: Model Development and Validation. *Journal of medical Internet research* 2020, 22:e24018.
- [20] Alimadadi A, Aryal S, Manandhar I, Munroe PB, Joe B, Cheng X: Artificial intelligence and machine learning to fight COVID-19. American Physiological Society Bethesda, MD, 2020.
- [21] Zhou Y, Wang F, Tang J, Nussinov R, Cheng F: Artificial intelligence in COVID-19 drug repurposing. *The Lancet Digital Health* 2020.
- [22] Landi I, Glicksberg BS, Lee HC, Cherng S, Landi G, Danieletto M, Dudley JT, Furlanello C, Miotto R: Deep representation learning of electronic health records to unlock patient stratification at scale. *NPJ Digit Med* 2020, 3:96.
- [23] Choi E, Bahadori MT, Schuetz A, Stewart WF, Sun J: Doctor ai: Predicting clinical events via recurrent neural networks. *Machine Learning for Healthcare Conference*, 2016. pp. 301-18.
- [24] Miotto R, Li L, Kidd BA, Dudley JT: Deep patient: an unsupervised representation to predict the future of patients from the electronic health records. *Scientific reports* 2016, 6:1-10.
- [25] Shickel B, Tighe PJ, Bihorac A, Rashidi P: Deep EHR: a survey of recent advances in deep learning techniques for electronic health record (EHR) analysis. *IEEE journal of biomedical and health informatics* 2017, 22:1589-604.
- [26] Choi E, Bahadori MT, Song L, Stewart WF, Sun J: GRAM: graph-based attention model for healthcare representation learning. *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2017. pp. 787-95.
- [27] Choi E, Xu Z, Li Y, Dusenberry M, Flores G, Xue E, Dai A: Learning the Graphical Structure of Electronic Health Records with Graph Convolutional Transformer. *Proceedings of the AAAI Conference on Artificial Intelligence*, 2020. pp. 606-13.
- [28] Wanyan T, Kang M, Badgeley MA, Johnson KW, De Freitas JK, Chaudhry FF, Vaid A, Zhao S, Miotto R, Nadkarni GN: Heterogeneous Graph Embeddings of Electronic Health Records Improve Critical Care Disease Predictions. *International Conference on Artificial Intelligence in Medicine: Springer*, 2020. pp. 14-25.
- [29] Veličković P, Cucurull G, Casanova A, Romero A, Lio P, Bengio Y: Graph attention networks. *arXiv preprint arXiv:171010903* 2017.
- [30] Wang X, Ji H, Shi C, Wang B, Ye Y, Cui P, Yu PS: Heterogeneous graph attention network. *The World Wide Web Conference*, 2019. pp. 2022-32.
- [31] Choi E, Bahadori MT, Sun J, Kulas J, Schuetz A, Stewart W: Retain: An interpretable predictive model for healthcare using reverse time attention mechanism. *Advances in neural information processing systems* 2016, 29:3504-12.
- [32] Barish M, Bolourani S, Lau LF, Shah S, Zanos TP: External validation demonstrates limited clinical utility of the interpretable mortality prediction model for patients with COVID-19. *Nature Machine Intelligence* 2020:1-3.
- [33] Gupta RK, Marks M, Samuels TH, Luintel A, Rampling T, Chowdhury H, Quartagno M, Nair A, Lipman M, Abubakar I: Systematic evaluation and external validation of 22 prognostic models among hospitalised adults with COVID-19: An observational cohort study. *European Respiratory Journal* 2020, 56.

- [34] D'Amour A, Heller K, Moldovan D, Adlam B, Alipanahi B, Beutel A, Chen C, Deaton J, Eisenstein J, Hoffman MD: Underspecification presents challenges for credibility in modern machine learning. arXiv preprint arXiv:201103395 2020.
- [35] Cyganek B, Graña M, Krawczyk B, Kasprzak A, Porwik P, Walkowiak K, Woźniak M: A survey of big data issues in electronic health record analysis. *Applied Artificial Intelligence* 2016, 30:497-520.
- [36] Xu J, Glicksberg BS, Su C, Walker P, Bian J, Wang F: Federated Learning for Healthcare Informatics. *J Healthc Inform Res* 2020:1-19.
- [37] Khosla P, Teterwak P, Wang C, Sarna A, Tian Y, Isola P, Maschinot A, Liu C, Krishnan D: Supervised Contrastive Learning. 2020. p. arXiv:2004.11362.
- [38] Xu J, Xu Z, Walker P, Wang F: Federated patient hashing. *Proceedings of the AAAI Conference on Artificial Intelligence*, 2020. pp. 6486-93.
- [39] Xu J, Xu Z, Yu B, Wang F: Order-Preserving Metric Learning for Mining Multivariate Time Series. 2020 IEEE International Conference on Data Mining (ICDM): IEEE, 2020. pp. 711-20.
- [40] Zhang X, He L, Chen K, Luo Y, Zhou J, Wang F: Multi-view graph convolutional network and its applications on neuroimage analysis for parkinson's disease. *AMIA Annual Symposium Proceedings: American Medical Informatics Association*, 2018. p. 1147.
- [41] Zhang X, Chou J, Wang F: Integrative analysis of patient health records and neuroimages via memory-based graph convolutional network. 2018 IEEE International Conference on Data Mining (ICDM): IEEE, 2018. pp. 767-76.
- [42] Li Z, Roberts K, Jiang X, Long Q: Distributed learning from multiple EHR databases: Contextual embedding models for medical events. *J Biomed Inform* 2019, 92:103138.
- [43] Zhang Y, Jiang H, Miura Y, Manning CD, Langlotz CP: Contrastive Learning of Medical Visual Representations from Paired Images and Text. 2020. p. arXiv:2010.00747.
- [44] Bordes A, Usunier N, Garcia-Durán A, Weston J, Yakhnenko O: Translating embeddings for modeling multi-relational data. *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 2*. Lake Tahoe, Nevada: Curran Associates Inc., 2013. pp. 2787-95.
- [45] Lipton ZC, Kale DC, Wetzell R: Modeling missing data in clinical time series with rnns. *Machine Learning for Healthcare* 2016, 56.
- [46] Levy O, Goldberg Y: Neural word embedding as implicit matrix factorization. *Advances in neural information processing systems* 2014, 27:2177-85.
- [47] Rajkomar A, Dean J, Kohane I: Machine learning in medicine. *New England Journal of Medicine* 2019, 380:1347-58.
- [48] Saito T, Rehmsmeier M: The precision-recall plot is more informative than the ROC plot when evaluating binary classifiers on imbalanced datasets. *PloS one* 2015, 10:e0118432.

A Architecture with Contrastive Loss


B Representation Space

u : patient of interest
 j : random patient

C_p : patient embedding
 C_e : event embedding


*: opposite of u 's label

contrastive loss shown in Eq. 7




similarity between patients p maximized by loss L_{pp} or L_{pp}^*
 similarity between patient p and event e maximized by loss L_{ep} or L_{ep}^*
 outcome relation


CL pushes positive patient/events away from negative ones



C Time Binning

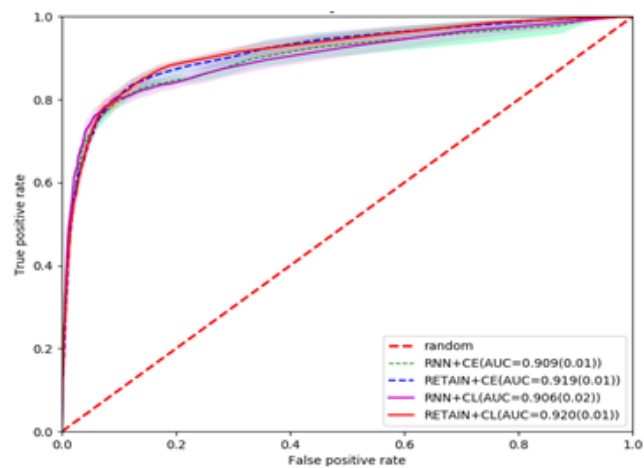


D Selection of Event Timing

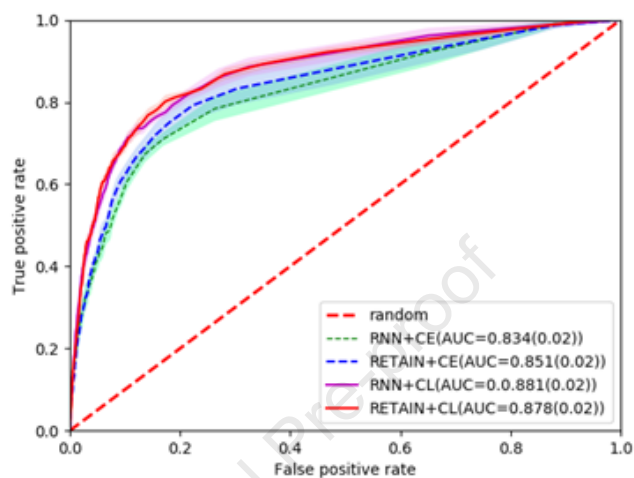


Mortality


Full Sample




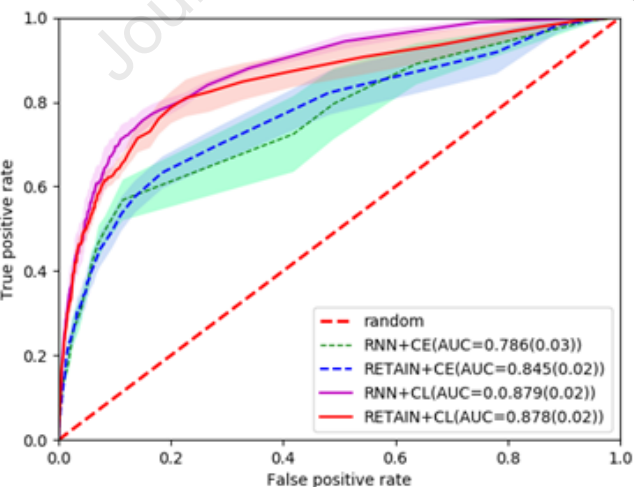

Intubation



ICU Transfer



Restricted Sample

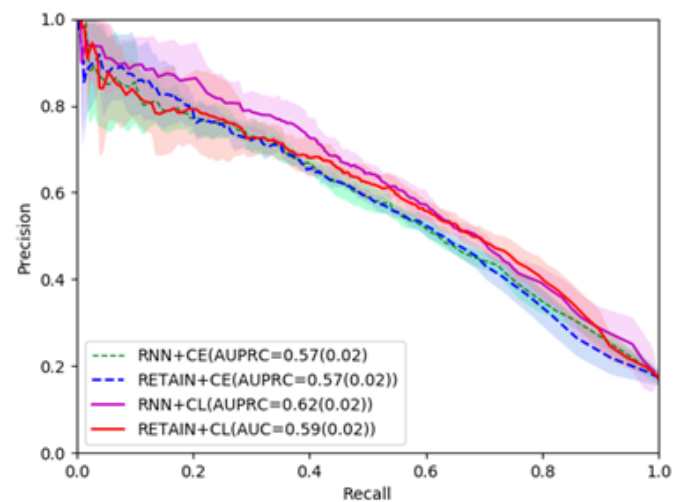

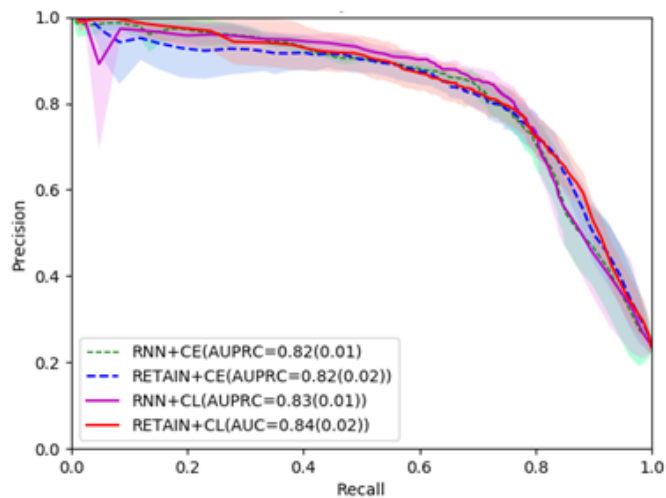


Mortality


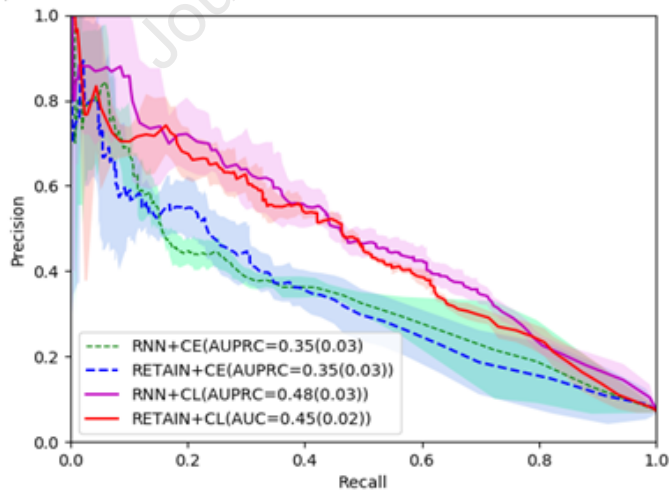
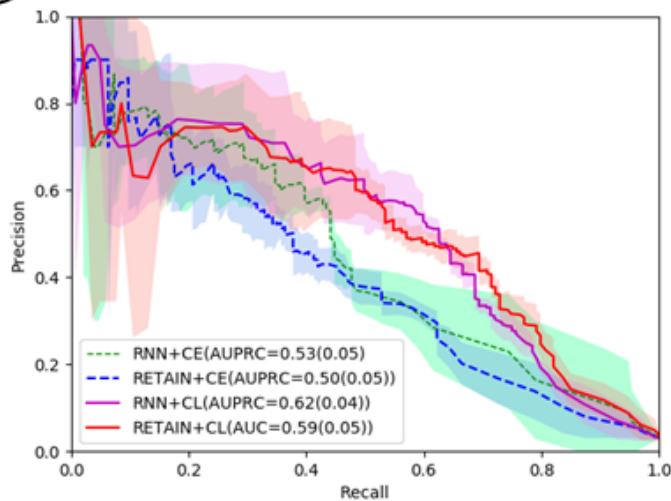
Intubation

ICU Transfer

Full Sample



Restricted Sample







Mortality

Contrastive Loss





Cross-Entropy Loss

Contrastive Loss





Cross-Entropy Loss



Intubation



ICU Transfer



(A) Contrastive Loss (B) Cross-Entropy Loss (C) Contrastive Loss (D) Cross-Entropy Loss

	Full Sample				Restricted Sample											
	6	12	18	24	6	12	18	24								
Temperature	0	0	0	0.01	0	0	0	0	0	0	0	0	0	0		
Respiration rate	0	0	0	-0.01	0	0	0	-0.01	0.01	0	0	0	0	-0.01		
Heart rate	0.01	0.01	0.01	0	0.01	0.01	0.01	0.01	0.01	0.02	0.02	0.02	0.01	-0.01		
Pulse oximetry	0.02	0.03	0.06	0.21	0.02	0.03	0.08	0.26	0.03	0.05	0.08	0.35	0.02	0.02	0.05	0.09
Oxygen saturation	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Systolic blood pressure	0	0.01	0.01	0.03	0	0.01	0.01	0.03	0.01	0.02	0.01	0.05	0	0	0.01	0.01
Diastolic blood pressure	0.01	0.01	0.02	0.04	0.01	0.01	0.02	0.06	0.01	0.02	0.02	0.07	0.01	0	0.01	0.01
Alanine Transaminase (ALT)	0.06	0.04	0.14	0.02	0.07	0.04	0.13	0.03	0.07	0	0.07	0	0.01	0.04	0.13	0.01
Albumin	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.02	0.04	0	0	0	0.01
Alkaline Phosphatase (ALP)	0	0	0	0	0	0	0	0	0.01	0	0.01	-0.01	0	0.01	0	0
Anion gap	0.03	0.03	0.04	0.06	0.02	0.02	0.03	0.05	0.03	0.02	0.04	0.06	0.01	0.01	0.01	0.01
Aspartate Aminotransferase (AST)	0.23	0.27	0.72	0.07	0.25	0.3	0.69	0.12	0.28	0.07	0.55	0.08	0.08	0.38	0.78	0.06
Atypical lymphocyte percentage	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
B-Type Natriuretic Peptide (BNP)	-0.01	0	0	0	0	0	0	0	0	0	0.01	0	0	0	0	0
Band count	0	0	0	0.01	0	0	0	0.01	0	0	0	0	0	0.01	0	0
Bicarbonate	0	0	0	0.01	0	0	0	0.01	0	0	0	0.01	0	0	0	0
Blood Urea Nitrogen (BUN)	0.09	0.09	0.1	0.12	0.07	0.08	0.1	0.12	0.13	0.13	0.16	0.15	0.02	0.02	0.02	0
C Reactive Protein (CRP)	0.02	0.02	0.03	0.04	0.01	0.02	0.02	0.03	0.02	0.01	0.02	0.02	0	0	0	0.01
Calcium	0	0	0	0	0	0	0	0	0.01	0.01	0	0.01	0	0	0	0
Chloride	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Creatine Kinase-MB (CK-MB)	-0.01	0	0	0	0.01	0	0	0	0	0	0	0	0	0	0	0
Creatine Phosphokinase (CPK)	0.04	0.08	0.04	0	0.04	0.07	0.04	0	-0.01	0.13	0.12	0	0.01	0.16	0.02	0
Creatinine	0.04	0.03	0.03	0.04	0.04	0.04	0.05	0.06	0.07	0.06	0.07	0.07	0.02	0.02	0.01	0
D-Dimer	0.02	0.02	0.03	0.03	0.01	0.01	0.03	0.03	0.03	0.02	0.04	0	0.01	0.01	0	0
Direct bilirubin	0.02	0.02	0.01	-0.01	0.02	0.02	0.01	-0.01	-0.03	-0.01	0	-0.02	0	0	0	-0.01
Ferritin	0.02	0.02	0.04	0.02	0.02	0.02	0.07	0.02	0.01	0.02	0.05	0	0.01	0	0.02	-0.01
Fibrinogen	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Glucose	0.02	0.02	0.02	0.03	0.01	0.02	0.02	0.03	0.01	0.01	0.02	0.03	0.01	0.01	0	0.01
Hematocrit	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Hemoglobin	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Interleukin-6	0	0	0.01	0	0	0	0	0	0	0	0	0	0	0	0	0
International normalised ratio (INR)	0.01	0.01	0	0.01	0.01	0.01	0	0.01	0.02	0.01	0	0.01	0.01	0.01	0	0
Ionized calcium	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Iron	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Lactate	0.03	0.11	0.11	0.23	0.03	0.1	0.11	0.24	0.05	0.17	0.12	0.36	0.01	0.09	0.04	0.05
Lactate Dehydrogenase (LDH)	0.06	0.12	0.3	0.1	0.05	0.1	0.28	0.08	0.02	0.03	0.11	0.22	0.03	0.14	0.37	0.05
Lymphocyte count	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Lymphocyte percentage	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.02	0	0.02	0.01	0	0	0	0
Mean Corpuscular Hemoglobin Concentration (MCHC)	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Monocyte count	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Monocyte percentage	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.02	0.01	0.02	0.02	0	0.01	0	0.01
Neutrophil count	0.02	0.02	0.03	0.03	0.02	0.02	0.04	0.03	0.04	0.03	0.08	0.03	0.01	0.01	0.01	0.01
Neutrophil percentage	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Partial Thromboplastin Time (PTT)	0	0	0	0.01	0	0	0	0.01	0.01	0.01	0	0	0	0	0	0
Partial pressure of carbon dioxide (PCO2)	0	0	0.01	0.01	0	0	0.01	0.01	0.03	0.01	0	0	0.01	0.01	0.01	0
Partial pressure of oxygen (PO2)	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Platelet count	0	0.01	0	0	0	0.01	0	0	0	0.01	0.01	0	0	0	0	0
Potassium	0.01	0.02	0.02	0.02	0.01	0.01	0.02	0.02	0.02	0.02	0.02	0.03	0	0.01	0.01	0
Prothrombin time (PT)	0.01	0.01	0	0.01	0.01	0.01	0	0.01	0.01	0.01	0	0.01	0	0	0	0.01
Red Blood Cell Count	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Red blood cell distribution width (RDW)	0	0	0	0	0	0	0	0	0	0	0	-0.01	0	0	0	0
Serum protein	0	0	0	0	0	0	0	0	0	0	0	0.02	0	0	0	0
Sodium	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Total bilirubin	0.01	0.01	0.01	0	0.01	0.01	0.02	0	0	0	0	-0.02	0	0.01	0.01	0
Total iron binding capacity (TIBC)	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Transferrin saturation	0	0	0	0	-0.01	0	0	0	0	0	0	0	0	0	0	0
Troponin	0.01	0	0.01	0.02	0	0	0.01	0.01	0.08	0.01	0.09	0.14	0	0	0	0.02
White Blood Cell Count	0.03	0.03	0.03	0.04	0.02	0.03	0.03	0.04	0.06	0.05	0.07	0.06	0.01	0.03	0.02	0.02
Mean Corpuscular Volume (MCV)	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Mean Platelet Volume (MPV)	0	0	0	0	0	0	0	0	0	0	0.01	0	0	0	0	0
pH	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

1.0

0.8

0.6

0.4

0.2

0.0

Highlights

- Contrastive loss consistently improves performance compared to cross-entropy loss
- For imbalanced outcome data, only contrastive loss model maintains proper clustering
- Contrastive loss better identifies relevant features

In Brief

Deep learning models applied on EHR data often utilize cross-entropy loss (CEL) as the primary optimization function. But CEL may not be suitable for real-world scenarios with imbalanced data. We develop a learning framework that incorporates both CEL and contrastive loss (CL) to tackle this issue. Our framework achieves better predictive performance and feature interpretability, particularly for imbalanced data.

Bigger Picture

The abundance of health data provides exceptional opportunities for machine learning analyses to improve care in terms of enhanced screening, diagnosis, and prognosis. One such data type is electronic health records, which generally consist of demographics, diagnoses, laboratory tests, vital signs, medications, and clinical notes. While deep learning has emerged as a powerful analysis tool to process large-scale data by extracting useful patterns, creating robust and generalizable models with such data is notoriously challenging due to scale, complexity, and outcome imbalance. In this work, we develop and refine a new model architecture based on the recently proposed contrastive deep learning. As a relevant use case, we demonstrate the power of this framework for predicting critical events in COVID-19 patients as well as an enhanced ability to identify important features. Our work shows promise for datasets with high missingness and outcome imbalance that normally hinders performance.