

Survival Seq2Seq: A Survival Model based on Sequence to Sequence Architecture

Ebrahim Pourjafari

*Ortho Biomed Inc.
Toronto, ON*

EBRAHIM.POURJAFARI@ORTHOBIOMED.CA

Navid Ziaei

*Ortho Biomed Inc.
Toronto, ON*

NAVID.ZIAEI@ORTHOBIOMED.CA

Mohammad R. Rezaei

*University of Toronto
Toronto, ON*

MR.REZAEI@MAIL.UTORONTO.CA

Amir Sameizadeh

*Ortho Biomed Inc.
Toronto, ON*

AMIR.SAMEIZADEH@ORTHOBIOMED.CA

Mohammad Shafiee

*University Health Network
Toronto, ON*

MOHAMMAD.SHAFIEE@UHN.CA

Mohammad Alavinia

*University Health Network
Toronto, ON*

MOHAMMAD.ALAVINIA@UHN.CA

Mansour Abolghasemian

*Ortho Biomed Inc.
Toronto, ON*

MANSOUR.ABOLGHASEMIAN@ORTHOBIOMED.CA

Nick Sajadi

*Ortho Biomed Inc.
Toronto, ON*

NICK.SAJADI@ORTHOBIOMED.CA

Editor: Editor's name

Abstract

This paper introduces a novel non-parametric deep model for estimating time-to-event (survival analysis) in presence of censored data and competing risks. The model is designed based on the sequence-to-sequence (Seq2Seq) architecture, therefore we name it Survival Seq2Seq. The first recurrent neural network (RNN) layer of the encoder of our model is

made up of Gated Recurrent Unit with Decay (GRU-D) cells. These cells have the ability to effectively impute not-missing-at-random values of longitudinal datasets with very high missing rates, such as electronic health records (EHRs). The decoder of Survival Seq2Seq generates a probability distribution function (PDF) for each competing risk without assuming any prior distribution for the risks. Taking advantage of RNN cells, the decoder is able to generate smooth and virtually spike-free PDFs. This is beyond the capability of existing non-parametric deep models for survival analysis. Training results on synthetic and medical datasets prove that Survival Seq2Seq surpasses other existing deep survival models in terms of the accuracy of predictions and the quality of generated PDFs.

1. Introduction

The objective of survival analysis is to estimate the hitting time of one or more events in the presence of censored data [Cox and Oakes \(2018\)](#). In healthcare, such events can include time of death, inception of a disease, time of organ failure [Dekker et al. \(2008\)](#), etc. The ability to predict the time or probability of certain events happening to a patient is a valuable asset for medical professionals and clinical decision makers, as it enables them to manage clinical resources more efficiently, make better informed decisions on the treatment they offer to patients, find the best organ donor-recipient matches, etc. Electronic health records (EHRs) often have characteristics that pose several challenges to the development of reliable survival models. Some EHRs are longitudinal, i.e., multiple observations per covariate per patient over time are recorded [Barak-Corren et al. \(2017\)](#). The survival model must be able to process such measurements and learn from their sequential temporal trends. Medical records, especially longitudinal observations, tend to be highly sparse. Therefore, any reliable survival model must effectively handle missing values, even if the missing rate is extremely high [Wells et al. \(2013\)](#); [Nazabal et al. \(2020\)](#). Dealing with right censored data is another complication of survival analysis models [Khan and Zubek \(2008\)](#). The right censored data happens often when medical centers lose track of a patient after a certain time, called the censoring time. Survival models must take into account censored data during the training phase. In addition, the presence of competing risk events is another challenge that survival models need to deal with. Having a long-tail distribution is another characteristic of many medical datasets [Rahman and Davis \(2013\)](#). In such skewed datasets, samples with shorter event times form the bulk of the distribution, while samples with longer event times make only a small portion of the dataset. Following the distribution of such skewed datasets is another challenge for survival and regression models. In fact, many existing survival models cannot accurately predict PDFs with a long time span from medical datasets.

The major challenge of developing survival analysis models is, however, the non-existence of the ground truth for the probability distribution of risk events [Clark et al. \(2003\)](#); [Bewick et al. \(2004\)](#). Medical records contain the time and type of events that happened to patients. However, the underlying distribution of time-to-event is unknown. This makes developing statistical models or supervised training of machine learning models for survival analysis difficult. Survival models can be divided into two main categories: parametric and non-parametric. Parametric survival models assume a certain stochastic distribution for the dataset, while trying to estimate the parameters of the assumed distribution. On the other hand, non-parametric survival models do not assume any prior distribution for the events.

Instead, they try to estimate distributions purely based on the observed relationship between the covariates and the time of events in the dataset.

To address the aforementioned challenges of developing survival analysis models and to alleviate the shortcomings of the existing survival models, we propose Survival sequence-to-sequence (Seq2Seq). Survival Seq2Seq is a non-parametric multi-event deep model, capable of processing longitudinal measurements with very high missing rates. The accuracy of our model in predicting event times as well as the quality of its generated probability distribution functions (PDFs) exceeds that of existing survival models. In addition, Survival Seq2Seq performs superbly on skewed datasets. The superiority of our model is backed by the results obtained by training Survival Seq2Seq on synthetic and medical datasets. These results will be provided in the later sections of this paper.

Our proposed Survival Seq2Seq model has the following key features:

- The first layer of the recurrent neural network (RNN)-based encoder network of Survival Seq2Seq is made of Gated Recurrent Units with Decay (GRU-D) [Che et al. \(2018\)](#) cells. GRU-D cells offer superior performance in imputing not-missing-at-random values. Taking advantage of GRU-D, Survival Seq2Seq can effectively handle high missing rates that commonly occur among medical datasets.
- The decoder network of our model is a recurrent network, which can generate substantially smoother PDFs compared to other non-parametric survival models. Since Survival Seq2Seq has fewer trainable parameters compared to a decoder made of dense layers, it suffers less from overfitting compared to other non-parametric models that use Multi-Layer Perceptron (MLP) in their decoders.
- We have enhanced the typical loss function used for training non-parametric survival models by improving the ranking loss term of our model. The improved ranking loss will help the model to better rank samples with longer event times.
- Our proposed model can be effectively trained on datasets with a long-tailed distribution, which is a common characteristic of healthcare datasets. This means that Survival Seq2Seq can accurately predict longer event times as well as the shorter ones.

Generalizable Insights about Deep Survival Models in the Context of Healthcare

The goal of survival analysis is to provide a hazard or reversely a survival function about a medical event for a patient, given some clinical observations. The hazard function or equivalently the hazard rate shows the probability of the occurrence of a medical event such as death or organ failure over time. While survival analysis has a long history in healthcare and non-healthcare applications using Cox-based proportional hazard (CPH) models, recent publications in the literature show the superior performance of deep learning models for predicting hazard functions. This means that the occurrence rate of medical events can be predicted more accurately over time. Consequently, healthcare providers can take preemptive actions more accurately and adequately.

As an example, the time of death of potential donors can be predicted via a hazard function, which allows procurement teams to make a timely attempt. Moreover, a deployed

tool based on an accurate hazard function can improve matchmaking and the outcome of transplants, leading to shorter wait lists and waiting times, improved longevity of the organs after transplantation, less need for re-transplantation, and longer survival of recipients with a higher quality of life. Currently, medical decision-making tools based on statistics studies of different cohorts and not fully supported by solid emerging evidence. For example, predicting mortality in ICU patients is of great value for assessing the severity of illness, adjusting the treatments to improve the patients’ chance of survival, adjudicating the impact of novel interventions, and resource management policies of healthcare systems Pirracchio et al. (2015); Yun et al. (2021). As our results show, our novel survival model significantly improves the accuracy and quality of mortality prediction for ICU patients. This enhancement in mortality prediction can improve the quality of decision making in all healthcare operations that rely on mortality prediction. In addition to promising results, once our model is deployed, the model can adapt to new patient data as it accumulates.

The rest of this paper is organized as follows: A short discussion of related works and their limitation is provided in section 2. The architecture of Survival Seq2Seq is described in Section 3. Section 4 represents the datasets used for training the model, while the experimental results of the training are provided in Section 5. Finally, section 6 concludes the paper.

2. Related Work

Classical parametric models rely on strong assumptions about the time-to-event distribution. Such strong assumptions allow these models to estimate the underlying stochastic process based on the observed relationship between covariates and time-to-event. However, the predicted probability distribution of these models is over-simplified and often unrealistic. CPH is an example of a parametric statistical model that simplifies the underlying distribution by assuming that the proportional hazard increases constantly over time. The model estimates the hazard function, $\lambda(t|\mathbf{x})$, the probability that an individual will experience an event within the time t , given all \mathbf{x} features (covariates). Although several works such as Vinzamuri et al. (2014); Vinzamuri and Reddy (2013); Li et al. (2016) have tried to address the shortcomings of the CPH model to some degree, the over-simplification of the underlying stochastic process limits the flexibility, generalizability, and the prediction power of the CPH model. Besides, the CPH model cannot process longitudinal measurements. Deep Survival Machines (DSM) Nagpal et al. (2021a) is an example of a machine learning-based parametric model that assumes a combination of multiple Weibull and Log-Normal primitive distributions for the survival function. A deep MLP model is trained to estimate the parameters of those distributions and a scaling factor to determine the weight of each distribution in the overall estimated hazard PDF. Training DSM is rather difficult, as the optimizer is easily diverged when trying to minimize its loss function and the model becomes overfit. Despite delivering a high Concordance Index (CI) score, the model performs poorly when estimating the hitting time of events. The CI score is a metric used for evaluating the ranking performance of survival models. The model is also unable to process longitudinal measurements.

The random survival forests (RSFs) method [Ishwaran et al. \(2008\)](#) is an extension of random forests that supports the analysis of right-censored data. The training procedure of RSFs is similar to other random forests. However, the branching rule is modified to account for right-censored data by measuring the survival difference between the samples on either side of the split, i.e., survival times. RSFs have become popular as a non-parametric alternative to CPH due to their less restrictive model assumptions. However, similar to CPH, RSFs cannot be trained on longitudinal datasets.

In [Katzman et al. \(2018\)](#), a non-parametric variation of CPH called DeepSurv is proposed as a treatment recommendation system. DeepSurv uses an MLP network for characterizing the effects of a patient’s covariates on their hazard rate. This solves the parametric assumptions of the hazard function in the original CPH model. It also leads to more flexibility of DeepSurv compared to CPH. As a result, DeepSurv outperforms other CPH-based models and it can learn complex relationships between an individual’s covariates and the effect of a treatment. DeepCox [Nagpal et al. \(2021b\)](#) proposes Deep Cox Mixtures (DCMs) for survival analysis, which generalizes the proportional hazards assumption via a mixture model, by assuming that there are latent groups and within each, the proportional hazards assumption holds. DCM allows the hazard ratio in each latent group, as well as the latent group membership, to be flexibly modeled by a deep neural network. However, both DeepSurv and DeepCox models still suffer from the same strong assumption of proportional hazards as the original CPH formulation. Also, neither of those two models supports longitudinal measurements.

As an alternative to assuming a specific form for the underlying stochastic process, [Lee et al. \(2018\)](#) proposes a non-parametric deep model called DeepHit, to model the survival functions for the competing risk events. Since no assumptions are made on the survival distribution, the relationship between covariates and the event(s) can now change over time. This is considered an advantage of DeepHit over CPH-based models. The first part of the DeepHit model, i.e., the encoder, is made of a joint MLP block. The decoder of the model is made of MLP blocks, each specific to one event. The output of each case-specific block is a discrete hazard PDF for each event. The last layer of each case-specific block contains N units, where each unit generates the likelihood for one timestep (*e.g.* one hour or one month) of the hazard PDF over the prediction horizon. A major drawback of this method for generating probability distributions is that the predictions of the output layer could arbitrarily vary from one unit to the next. This causes the overall generated PDFs to be extremely noisy. Moreover, depending on the number of time steps in the prediction horizon, the number of trainable parameters in the output layer could become too high and may cause overfitting when training the model.

Dynamic-DeepHit (DDH) [Lee et al. \(2019\)](#) is an extension of DeepHit, capable of processing longitudinal measurements. In this model, the encoder is replaced with RNN layers followed by an attention mechanism. The RNN block can learn the underlying relationship between longitudinal measurements and provide finer predictions compared to the MLP block in DeepHit. The case-specific blocks of DeepHit and DDH are the same, which means that DDH suffers from predicting noisy PDFs and overfitting. Recurrent networks are also utilized in [Ren et al. \(2019\)](#) to generate probability distributions of events. In contrast to non-parametric deep models like [Lee et al. \(2018\)](#) or [Lee et al. \(2019\)](#) that use MLP-based decoders, the authors in [Ren et al. \(2019\)](#) have utilized an RNN-based decoder to estimate

probability distributions of events. They have achieved higher CI scores compared to Lee et al. (2018), which can be contributed to the higher quality of RNN-generated probability distributions and less overfitting issues. Although RNN based, their survival model is not designed to process longitudinal measurements.

Recently, the application of the transformer architecture in survival analysis has been studied in some research such as Wang and Sun (2021) and Hu et al. (2021). In addition, studies like Engelhard et al. (2020) suggest distinguishing between the probability distribution of events and their hitting times due to the lack of ground truth for the probability distribution of events. Kamran and Wiens (2021) focuses on improving the calibration capability of survival models, while maintaining the discriminative power of these models. They argue that the calibration power of survival models is even more important than their ranking power. It is worth mentioning that none of the aforementioned studies considers longitudinal measurements in the modeling structure.

3. Methods: Survival Seq2Seq

A successful non-parametric survival model must accept longitudinal EHRs as the input, and must output hazard PDFs for competing risks, while using survival times as the ground truth for training. We realized that the Seq2Seq architecture had a strong potential for performing such a survival analysis. This architecture is commonly used for natural language processing (NLP) and time series decoding tasks Rezaei et al. (2022). For example, it can translate one language to another. A typical Seq2Seq model is made up of an RNN-based encoder network and another RNN-based network as the decoder. In case of language translation, the encoder encodes a sentence of the source language and sends the encoded sequence to the decoder. The decoder then generates the sentence of the destination language word by word based on the previously generated words in the destination sentence as well as the encoded sequence. Also, adding an attentional mechanism to a Seq2Seq model improves the performance of the model when translating longer sentences.

The Seq2Seq architecture can be adapted to perform survival analysis. In a nutshell, the encoder of a Seq2Seq model can process longitudinal EHR and send the encoded sequence to the decoder, while the decoder generates a discrete hazard PDF of an event based on the encoded sequence. Despite showing a strong potential, the original Seq2Seq architecture needs to undergo several modifications to be ready for performing survival analysis. The model needs a missing value handling mechanism to effectively impute the sheer number of missing values in longitudinal medical datasets. The decoder of Seq2Seq can generate a PDF for each competing risk. In addition, a proper loss function is required to train such a non-parametric model. While considering censored data, such a loss function must be able to shape the generated PDFs based on the observed relationship between measurements and event times.

Figure 1 shows the overall structure of Survival Seq2Seq, a survival model based on the Seq2Seq architecture. The model follows the basic Seq2Seq architecture with an encoder for processing longitudinal measurements and a decoder for generating hazard PDFs for multiple events. The first RNN layer of the encoder is made up of GRU-D cells. These cells are highly effective in imputing missing values of medical datasets. The decoder network is composed of several RNN-based decoder blocks, where each block is responsible for gener-

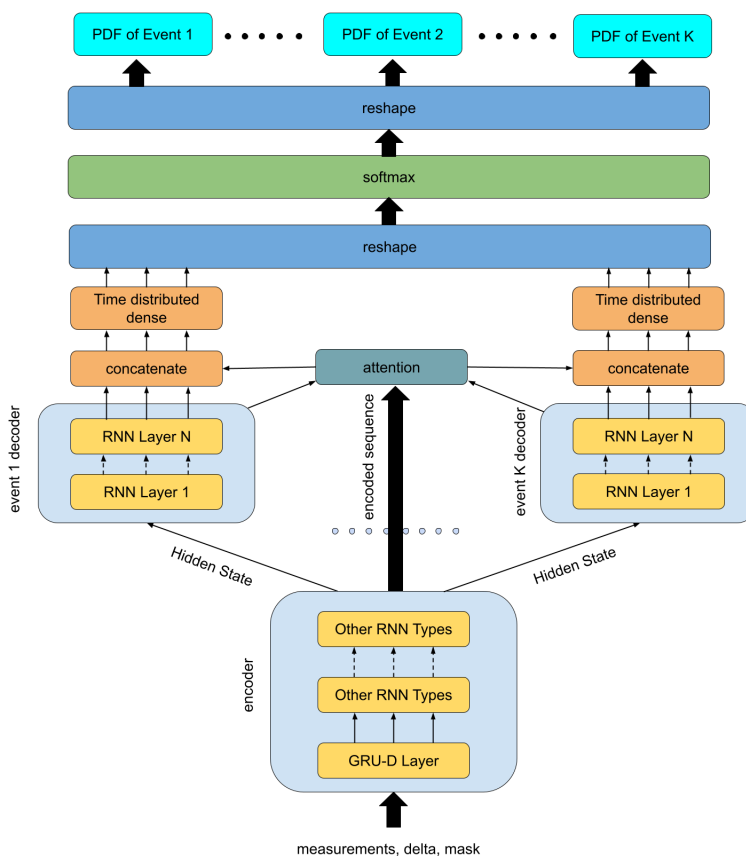


Figure 1: The overall architecture of Survival Seq2Seq

ating the PDF for one event in the dataset. We have added an attentional mechanism to Survival Seq2Seq to improve its overall performance. To train Survival Seq2Seq, we use a multi-term loss function composed of the log-likelihood loss [Lee and Whitmore \(2006\)](#) plus an improved ranking loss term that enhances the ranking performance of the model. Design of Survival Seq2Seq is discussed in greater detail in the rest of this section.

3.1. Encoder

The encoder network is responsible for processing longitudinal measurements of patients and passing the encoded sequence to the decoder network. As mentioned earlier, longitudinal EHRs are very sparse. To handle the missing values of longitudinal measurements, we use GRU-D cells in the first layer of the encoder network. The patterns of missing values and their relationship with labels in medical datasets provide valuable information for supervised learning tasks [Che et al. \(2018\)](#). A GRU-D cell learns the relationship between labels and the missing pattern of covariates during a supervised learning process and utilizes the observed relationship to impute missing longitudinal measurements for continuous covariates. GRU-D imputes missing values of a covariate by applying a decay rate to the last measured value of that covariate. The influence of the last covariate measurement is reduced proportionally

over time if the covariate has not been measured. If a covariate is not measured for a long period of time, GRU-D imputes that missing value by relying more on the mean value of that covariate rather than its last measurement. The imputation process in GRU-D needs three input vectors: a vector containing longitudinal measurements, a mask vector which indicates if a measurement is available or missing (value one represents a covariate is measured and value zero represents the covariate is missing) and a Delta vector which represents the time difference between the current timestamp and the last timestamp at which a covariate for a patient was measured. Readers can refer to [Che et al. \(2018\)](#) for the implementation details of GRU-D cells. The encoder network of Survival Seq2Seq can be stacked with multiple RNN layers if desired. In that case, the stacked layers can be made up of some other recurrent cells such as Gated Recurrent Unit (GRU) or Long Short-Term Memory (LSTM). A stacked encoder allows the model to learn more complex temporal relationships in data, provided that the training set is large enough to support a deeper network. If not, the encoder can also work with only one GRU-D layer.

3.2. Decoder

As depicted in Figure 1, the decoder network is made of K decoder blocks, each specific to one event in the dataset. The censored event is not considered as an interested event, therefore no decoder block is assigned for censored records. Decoder blocks can be made of stacked RNN layers of vanilla RNN, LSTM or GRU. Similar to the encoder network, a stacked decoder block can be used to learn more complex relationships in data. The generated hazard PDFs are discrete, meaning that the prediction horizon is divided into several timesteps (bins) and a decoder block predicts the value (likelihood) of each bin sequentially. To generate the likelihood for a given timestep, the decoder relies on the likelihood of the previous timestep as well as the encoded sequence. Therefore, the value of the PDF for any timestep is dependent on the previous timesteps and cannot arbitrarily change. As a result, the generated PDFs are smooth and ripple-free.

The output of each decoder block is concatenated with the attention vector and then passed through a time distributed dense layer with a *relu* activation function. The attentional mechanism improves the performance of the model in dealing with long encoded sequences, i.e, data samples with too many measurements. The output tensor of all decoder blocks is reshaped to a one-dimensional tensor, where a *softmax* activation is applied on the resulted tensor. The *softmax* activation guarantees that the joint cumulative distribution function (CDF) of all events is always equal to one. The output of the *softmax* layer is then reshaped to its original dimensions. This provides the estimated hazard PDFs for all events as follows:

$$PDF = \begin{pmatrix} p_{1,1} & p_{1,2} & \cdots & p_{1,T_h} \\ p_{2,1} & p_{2,2} & \cdots & p_{2,T_h} \\ \vdots & \vdots & \ddots & \vdots \\ p_{K,1} & p_{K,2} & \cdots & p_{K,T_h} \end{pmatrix}, \quad (1)$$

where each row represents the predicted hazard PDF for an event, while T_h is the number of timesteps (length of decoder) for the estimated PDFs. Using Equation 1, the estimated hazard CDF for event k^* at timestep τ for a set of covariates x is given by

$$CDF_{k^*}(\tau|x) = \sum_{i=0}^{\tau} p_{k^*,i} \quad (2)$$

Depending on the use case, either of the PDF formulated in Equation 1 or the CDF in Equation 2 can be considered as the output of the model. Also, if the hitting time of an event is desired, the expected value of the PDF can be considered as the predicted time of the event.

3.3. Loss function

The loss function used for training Survival Seq2Seq is $L = L_l + L_r$ in which L_l and L_r are log-likelihood Lee and Whitmore (2006) and ranking terms, respectively. The log-likelihood term is defined as follows:

$$L_l = - \sum_{j \in U_{uc}} \log(p_{k_t, \tau_t}^j) - \sum_{j \in U_c} \log(1 - \sum_{k=1}^K CDF_k(\tau_t|x_j)), \quad (3)$$

in which U_{uc} and U_c are the sets of uncensored and censored patients, respectively. The index k_t is the ground truth for the first hitting event, while τ_t is the time of the event or censoring. The log-likelihood loss is the main loss term for training Survival Seq2Seq. The ground-truth for the probability distribution of events is unknown to us. The log-likelihood loss allows a non-parametric model such as Survival Seq2Seq to be trained for predicting probability distributions, while using the first hitting event time as the ground truth. The first term of this loss is used for training the model on the first hitting event and its corresponding time for uncensored patients, while the second term is used for training the model on censored data. The first term of L_l is designed to maximize the estimated hazard PDF for event k_t at τ_t , while minimizing it for other timesteps.

L_r is a ranking loss and improves the ranking capability of the model. The idea behind the ranking loss is that if an event happens for a given patient at a given time, the estimated CDF for that patient at the time of the event must be higher than patients who have not experienced the event yet. Therefore, this loss term increases the overall CI score of the model. According to Jing et al. (2019) the ranking loss can be defined as:

$$L_r = -\frac{1}{|U_a|} \sum_{(i,j) \in U_a} \Phi(CDF_k(\tau_{t_i}|x_i) - CDF_k(\tau_{t_j}|x_j)), \quad (4)$$

where τ_{t_i} is time of event for patient i and $\Phi(\cdot)$ is a differentiable convex function where we use the exponential function $\Phi(z) = \exp(z)$. The U_a term is a set of acceptable patient pairs. In an acceptable pair, the first element must be uncensored, and the event for the first element must occur before the event or censoring time of the second element. Even though this loss function improves the ranking power of survival models, studies like Tjandra et al. (2021) suggest that the ranking power of this loss can be further improved. To improve the discriminative power of the ranking loss, we modify Equation 4 and propose the following:

$$L_r = -\frac{1}{|U_a|} \sum_t \sum_{(i,j) \in U_a} \Phi(CDF_k(\tau_{t_i}|x_i) - CDF_k(\tau_{t_j}|x_j)). \quad (5)$$

The difference between Equation 5 and Equation 4 is that the updated ranking loss compares two elements of a given pair at every timestep over the prediction horizon, while the loss in Equation 4 only compares the pair at the time of the event. Our experimental results provided in the upcoming sections suggest that the updated ranking loss improves the ranking capability of the model, especially for events with longer hitting times. A longer hitting time means more time steps at which the modified ranking loss evaluates an acceptable pair. This translates to a better ranking performance on events with longer hitting times.

4. Cohort

The datasets used for training the model are described in the following:

SYNTHETIC: To evaluate the performance of Survival Seq2Seq, we created a synthetic dataset based on a statistical process. Here, we consider $\mathbf{x} = (x^0, \dots, x^K)$ as a tuple of K random variables in a way that each random variable has an isotropic Weibull distribution

$$f(x; \gamma, \mu, \alpha) = \frac{\gamma}{\alpha} \left(\frac{x - \mu}{\alpha} \right)^{(\gamma-1)} \exp(-((x - \mu)/\alpha)^\gamma) \quad x \geq \mu; \gamma, \alpha > 0,$$

where γ is the shape parameter, μ is the location parameter and α is the scale parameter. We model the distribution of two event times, $T_i^{(1)}$ and $T_i^{(2)}$, for each data sample i as a nonlinear combination of these K random variables at time index i defined by

$$T_i^{(1)} = f(\boldsymbol{\alpha}^T \times (\mathbf{x}_i^{k_1})^2 + \boldsymbol{\beta}^T \times (\mathbf{x}_i^{k_2})), \quad (6)$$

$$T_i^{(2)} = f(\boldsymbol{\alpha}^T \times (\mathbf{x}_i^{k_2})^2 + \boldsymbol{\beta}^T \times (\mathbf{x}_i^{k_1})). \quad k_1, k_2 \in \{1, \dots, K\} \quad (7)$$

where k_1 and k_2 are two randomly-selected subsets of K covariates that satisfy $k_1 \cap k_2 = \emptyset, k_1 \cup k_2 = K$. By selecting the Weibull distribution in our synthetic data generator, the event times will be exponentially distributed with an average that depends on a linear (with parameters set $\boldsymbol{\beta}$) and quadratic (with parameters set $\boldsymbol{\alpha}$) combination of the random variables. This means long-tailed distributions for the event times similar to medical datasets, as shown in Figure 2. The figure shows that the number of samples in the histogram are not monotonically decreasing over the time. Instead, the number of samples for each event peaks at some time and then decreases with a long tail. We chose this deliberately to evaluate the performance of the model on a dataset with a complex time-to-event distribution. For each data sample, we have $(\mathbf{x}_i, \delta_i, T_i)$, where $\delta_i = \min\{T_i^{(1)}, T_i^{(2)}\}$ identifies the type of the event happening at T_i , i.e., the event time. In sum, we considered $K = 20$ and generated 20000 data samples from the defined stochastic process with 20% censoring rate, and a not-missing-at-random rate of 77% and a maximum event time of 200. To simulate longitudinal measurements, we assume that the covariates for each sample are measured a random number of times, while the measured values at each time stamp are increased or decreased nonlinearly with a random rate specific to each sample-covariate pair.

MIMIC-IV: The MIMIC-IV dataset Johnson et al. (2020) contains hospital clinical records for patients admitted to a tertiary academic medical centre in Boston, MA, USA., between 2008-2019. This database contains demographics data, laboratory measurements, medications administered, vital signs and diagnosis of more than 71000 ICU patients. We




Figure 2: The histogram of survival times in the synthetic dataset. The histogram shows that the generated dataset is long-tailed similar to medical datasets.

considered a 33-day prediction horizon with a 4-hour time resolution for the decoder of Survival Seq2Seq. Therefore, patients with an event time longer than 33 days were considered censored. We also selected 108 covariates from this dataset based on the feedback of our medical team as well as by applying feature selection methods to the dataset.

5. Results

The performance of Survival Seq2Seq was evaluated by training the model on SYNTHETIC and MIMIC-IV datasets. Results were compared to DDH for benchmarking, as DDH was the only survival model known to us, capable of processing longitudinal measurements for survival analysis.

5.1. Evaluation Approach/Study Design

Mean absolute error (MAE) is the main metric we use for evaluating the performance of the models. MAE is defined as the mean of the absolute difference between the predicted time of an event and the observed time for that event for uncensored data samples. We consider the expected value of a predicted PDF as the predicted event time. The other metric used in this paper is the time-dependent CI score [Antolini et al. \(2005\)](#) defined as follows:

$$\mathbb{C}\mathbb{I}(t) = P(\hat{F}(t|x_i) > \hat{F}(t|x_j) | \delta_i = 1, T_i < T_j, T_i \leq t),$$

where $\hat{F}(t|x_i)$ is the estimated CDF of the event, truncated at time t , given a set of covariates x_i . This metric evaluates the performance of the models when predicting the order of events for data samples. The CI measure can evaluate the performance on both censored and uncensored samples in the test dataset. The time dependency of this metric allows us to evaluate the performance of models when capturing the possible changes in risk over time.

The results will be reported in four 25%, 50%, 75% and 100% quantiles with 5-fold cross-validation. We consider one GRU-D layer for the encoder and one GRU layer for each decoder block of the decoder network. For the Synthetic dataset, the encoder contains 100 GRU-D cells, while each decoder block (there are two decoder blocks, each for one event) is made of 100 GRU cells. To train the model on MIMIC-IV, we consider one GRU-D

layer for the encoder and a GRU layer for decoder, each containing 120 cells. We also applied 40% recurrent dropout rate to the RNN layers for training on both datasets. The optimal number of cells and dropout rate were obtained using grid search hyperparameter optimization. 50 epochs are considered for training the model on each fold with an early stopping condition in place.

In addition, we consider the length of the decoder of Survival Seq2Seq 25% longer than the maximum event time of each dataset. This is necessary for dealing with censored samples at the maximum event time (200 for SYNTHETIC and 33 days for MIMIC-IV) or a time close to the defined maximum event time. This extra decoder length allows the model to shift a considerable portion of the PDF for those samples to later timesteps, so that the predicted time of the event for censored data happens after the censoring time. The extra length of the decoder is a hyperparameter of the model, where its optimal value can be chosen based on the performance of the model on censored data; if the model performs poorly on censored data, the length of the decoder should be increased until satisfactory performance on censored data is achieved.

5.2. Evaluating Using the SYNTHETIC Dataset

Table 1 compares the performance of Survival Seq2Seq and DDH on the SYNTHETIC dataset in terms of MAE. The mean and variance of the results of the five folds are used for calculating the confidence interval of predictions. It can be seen that Survival Seq2Seq significantly outperforms DDH for both events in all quantiles except the first quantile. Figure 3 represents the difference between the ground truth and prediction for Survival Seq2Seq and DDH for a few uncensored samples of the SYNTHETIC dataset. It can be observed that Survival Seq2Seq can follow the distribution of the dataset by predicting values close to the ground truth, whether the survival time is closer to zero or it is much longer. On the other hand, DDH performs very poorly in predicting longer event times. Prediction results provided in Table 1 as well as Figure 3 proves that Survival Seq2Seq can effectively predict the event time on datasets with long-tailed distributions.

As Figure 3 reveals, DDH has a tendency to predict shorter event times for all events, whether the ground truth is long or short. This translates to a lower MAE for shorter event times for DDH compared to Survival Seq2Seq. This situation is somehow similar to an imbalanced binary classification in which a naively-trained classification model leans toward predicting in favor of the class with the higher share of training samples. The overall accuracy of such a classification model could be very high, although the model would perform very poorly on the minority class. Using the same analogy, a better way to evaluate a model like DDH with skewed predictions is to look at its performance in the last quantile, where events with longer hitting times can be found. This is where Survival Seq2Seq outperforms DDH.

The slightly higher MAE of Survival Seq2Seq compared to DDH in the first quantile can also be explained based on the imputation mechanism of GRU-D cells. The number of measurements for data samples with very short event times are scarce. If the value of a covariate for a given data sample is missing at the first measurement timestamp, GRU-D has to rely solely on the mean of that covariate to impute that missing value. This is

		Quantiles			
		25%	50%	75%	100%
Survival Seq2Seq	event 1	11.85±0.6	12.47 ± 1.2	14.01 ± 1.4	15.54 ± 1.8
	event 2	12.26 ± 0.5	13.20 ± 0.9	15.55 ± 2.2	20.83 ± 3.8
DDH	event 1	8.79 ± 0.7	20.36 ± 1.8	29.93 ± 1.8	35.03 ± 1.8
	event 2	10.32 ± 0.3	18.45 ± 1.4	33.96 ± 1.5	52.22 ± 1.2

Table 1: Comparison between the MAE of Survival Seq2Seq and DDH on the SYNTHETIC dataset. Results are reported with 95% confidence interval.

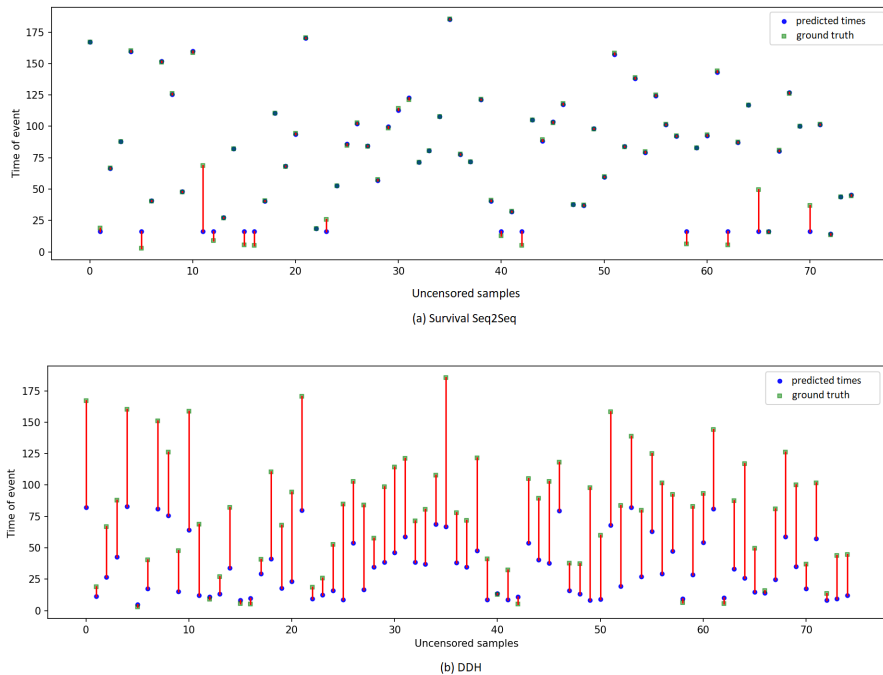


Figure 3: The difference between the ground truth (green squares) and predictions (blue dots) of Survival Seq2Seq (top) and DDH (bottom) on a few uncensored samples from the SYNTHETIC dataset. Red lines show the difference between predictions and ground truth for samples. The shorter the red line, the more accurate the prediction. The comparison shows that Survival Seq2Seq follows the distribution of the event times more accurately than DDH. The difference between the performance of the two models is more apparent for longer event times.

obviously not ideal and causes higher prediction errors for very early predictions. However, the imputation performance of GRU-D improves as more measurements are accumulated.

Table 2 represents the time-dependent CI score for the two models on the SYNTHETIC dataset. Despite providing marginally lower CI scores for the two events in the first quantile, our model outperforms DDH in all other quantiles, with a significant 0.08 higher average

		Quantiles			
		25%	50%	75%	100%
Survival Seq2Seq	event 1	0.835±0.04	0.774±0.06	0.827±0.02	0.844±0.02
	event 2	0.850±0.03	0.804±0.04	0.808±0.03	0.835±0.03
DDH	event 1	0.842±0.01	0.766±0.01	0.760±0.01	0.776±0.01
	event 2	0.853±0.01	0.787±0.01	0.744±0.01	0.743±0.01

Table 2: Comparison between the time-dependent CI score of Survival Seq2Seq and DDH on the SYNTHETIC dataset. Results are reported with 95% confidence interval.

		Quantiles			
		25%	50%	75%	100%
Survival Seq2Seq	event 1	0.835±0.04	0.774±0.06	0.827±0.02	0.844±0.02
	event 2	0.850±0.03	0.804±0.04	0.808±0.03	0.835±0.03
Survival Seq2Seq without ranking loss	event 1	0.813±0.09	0.729±0.10	0.751±0.06	0.779±0.07
	event 2	0.832±0.07	0.755±0.09	0.731±0.06	0.769±0.06

Table 3: Comparison between the CI scores of Survival Seq2Seq with the ranking loss 5 (same as the CI score of Survival Seq2Seq in Table 2) and Survival Seq2Seq without a ranking loss. The comparison suggests that the ranking loss has a strong effect on the ranking capability of the model across all quantiles.

CI score compared to DDH on the last quantile. The higher CI score of Survival Seq2Seq with respect to DDH can be contributed to the introduction of the new ranking loss in Equation 5. As the results suggest, the new ranking loss improves the ranking performance of Survival Seq2Seq compared to a model like DDH that employs a ranking loss similar to Equation 4 for improving its ranking performance.

To further investigate the impact of the proposed ranking loss 5 on the ranking capability of Survival Seq2Seq, we performed a study in which we turned this loss off and allowed the model to be trained solely using the log-likelihood 3 loss function. Table 3 presents the CI score of Survival Seq2Seq with and without the ranking loss. It is evident that turning the ranking loss off considerably reduces the CI score of the model across all quantiles.

5.3. Evaluating Using the MIMIC-IV Dataset

The performance of Survival Seq2Seq on MIMIC-IV is provided in Table 4. This table shows the MAE and CI scores for our model and compares them to the outcome of DDH. One can observe a pattern similar to Table 1 when comparing the MAE of Survival Seq2Seq and its counterpart. The MAE of our model is marginally higher than DDH in the first quantile. However, Survival Seq2Seq beats DDH in all other quantiles, so that the mean absolute error of our model is less than half of the MAE of DDH on 100% of the test data. Besides, the performance of Survival Seq2Seq in terms of the CI score exceeds the outcome of DDH for all quantiles. Again, the updated ranking loss in Equation 5 is a factor that contributes to these high CI scores.




Figure 4: The difference between the ground truth (green squares) and predictions (blue dots) of Survival Seq2Seq (top) and DDH (bottom) on a few uncensored samples from MIMIC-IV. Each red line indicates the difference between the prediction and the actual time of death for a patient. The shorter the red line, the more accurate the prediction.

Figure 4 plots the difference between the predicted time and actual time of death for a few randomly-selected deceased patients of MIMIC. The red vertical lines show the difference between the predicted time of death and ground truth. One can observe that the predictions of Survival Seq2Seq on MIMIC-IV are not as accurate as its predictions on the Synthetic dataset presented in Figure 3. This is not an unexpected phenomenon, as a real-world medical dataset like MIMIC-IV is influenced by many complicated dynamics that a synthetic dataset can never simulate. Figure 4 also shows that Survival Seq2Seq provides more precise predictions compared to the other model. The difference between the two models is more evident on patients with longer times of event.

The superior imputation performance of GRU-D cells compared to other imputation techniques in classification problems has been studied and shown in Che et al. (2018) on several longitudinal datasets including MIMIC. To study the imputation quality of GRU-D cells in the context of survival analysis, we created a modified version of Survival Seq2Seq that does not utilize GRU-D cells. Instead, we added a masking vector to the input of the model that indicates if a measurement is available by taking one or it is missing by taking zero. This masking technique is a common method for handling missing values of longitudinal datasets. All the GRU-D cells in this modified model were replaced with

		Quantiles			
		25%	50%	75%	100%
Survival Seq2Seq	MAE	34.83±4.1	37.06±4.6	39.53±4.0	62.74±3.2
	CI	0.876±0.02	0.882±0.02	0.885±0.02	0.906±0.02
DDH	MAE	32.88±8.4	40.23±2.4	61.17±2.8	125.95±4.7
	CI	0.863±0.03	0.849±0.02	0.845±0.03	0.528±0.04

Table 4: Comparison between the results of Survival Seq2Seq and DDH on MIMIC-IV. Results are reported with 95% confidence interval.

		Quantiles			
		25%	50%	75%	100%
Survival Seq2Seq		34.83±4.1	37.06±4.6	39.53±4.0	62.74±3.2
Survival Seq2Seq without GRU-D		43.82±5.7	47.28±5.1	49.71±4.8	65.57±4.0

Table 5: Comparison between the MAE of the original Survival Seq2Seq with GRU-D (same as the MAE of Survival Seq2Seq reported in Table 4) and the survival Seq2Seq with masking on MIMIC-IV. The comparison shows that the original model that takes advantage of GRU-D cells has lower prediction error across all quantiles.

GRU cells. Table 5 compares the MAE of this version of Survival Seq2Seq without GRU-D cells to the original Survival Seq2Seq on MIMIC-IV. The MAE for the original model is considerably lower than the MAE for Survival Seq2Seq without GRU-D imputation in all quantiles. Considering the fact that the only difference between the two models is their imputation approach, one can argue that GRU-D indeed offers a superior imputation performance compared to conventional methods of handling missing values in the context of survival analysis.

The quality of generated PDFs is another important factor besides the prediction accuracy when comparing Survival Seq2Seq to other non-parametric models. The superiority of our model becomes apparent from Figure 5, where the predicted PDFs of Survival Seq2Seq and DDH are compared for five randomly selected uncensored patients of MIMIC-IV. As described in earlier sections, the RNN-based decoder of our model can generate smooth and ripple-free probability distributions. This is in contrast to DDH where using an MLP-based decoder results in PDFs with high fluctuations, as shown in Figure 5.

6. Discussion

Each of the MAE and CI metrics used in this paper is useful for evaluating a different use case of a survival model. MAE is useful for evaluating the outcome of a survival model when predicting the time of events (calibration) is desired, while CI is a better choice when the model is applied on ranking problems. However, we cannot train a model that provides optimal MAE and CI at the same time. The log-likelihood loss has a higher influence on the calibration capability of the model than the ranking loss. Therefore, to minimize MAE, one must assign a higher weight to the log-likelihood loss than the ranking loss. On the




Figure 5: (a) The PDFs generated by Survival Seq2Seq for five randomly selected uncensored patients of MIMIC-IV and (b) The generated PDFs of DDH for the same patients. The ground truth for death and the predicted time of death of the two models for each patient is shown at the bottom of the figure. The comparison shows that the quality of the predicted probability distribution of Survival Seq2Seq is superior than DDH.

other hand, the ranking loss must be assigned a higher weight if the ranking capability of the model and a higher CI score is desired. Since we believe that calibration is a more important task than ranking, we assigned a higher weight to the log-likelihood loss to minimize MAE. We did the same for DDH to make the comparison between the two models fair. Assigning a higher weight to the ranking loss results in a maximized CI score. For example, we achieved a CI score of 0.932 ± 0.01 on the last quantile of MIMIC-IV with Survival Seq2Seq when a higher weight was assigned to the ranking loss. This is almost 3% higher than the CI score reported in Table 4 when calibration was the objective.

MAE is the metric of choice for calibration problems. This score is symmetric, *i.e.* it does not care about the sign of the error and whether the model is overpredicting or underpredicting. Depending on the context and nature of the survival problem, it might be useful to consider a calibration metric that evaluates a model differently for overshooting and undershooting. Such a metric helps to better evaluate survival models in situations that overpredicting or underpredicting is costlier than the alternative.

Survival Seq2Seq can accurately follow the distribution of the event time for long-tail datasets, while DDH and many other survival or regression models cannot. This is an interesting feature of our model and we mainly consider this as a side contribution of the RNN-based decoder network to our model. We were initially not sure if this feature of the model is a result of using GRU-D cells in the encoder of our model, the modified ranking loss, or the use of an RNN-based decoder. We created a second version of Survival Seq2Seq with the same RNN-based decoder network as the original model, but without the GRU-D layer in the encoder and without the modified ranking loss. We used a simple masking technique for dealing with missing measurements. This version of Survival Seq2Seq was still able to follow the distribution of the time-to-event on long-tail datasets, although the accuracy of the model dropped. Therefore, we concluded that the RNN-based decoder was the main contributor to this interesting feature of our model in performing well on long-tail datasets. However, the exact underlying mechanism of this phenomenon is unknown to us and requires further investigation.

Based on the provided results in terms of MAE and CI score, we proved that Survival Seq2Seq has a high prediction accuracy. However, accuracy metrics are not the only factor that must be taken into account when evaluating the performance of a survival model. We believe that the quality of predicted PDFs is as important as the accuracy metrics. As the provided PDF sample in this paper shows, the generated PDFs of Survival Seq2Seq are exceptionally smooth and ripple-free. This is an outstanding feature for a non-parametric survival model.

Limitations

The GRU-D layer that makes the first layer of the encoder of our model relies on previously measured values and the mean of covariates to impute missing data. If a covariate for a row of data (a patient) is missing while it is not previously measured, GRU-D has no choice but to impute that missing value with the mean of the covariate. This is not an ideal behavior and we believe that it is contributing to lower performance of Survival Seq2Seq in the first quantile of data, where there are not enough measurements for GRU-D to properly impute missing variables. We acknowledge this limitation of our model on early predictions and encourage the readers to think of ways that could mitigate this phenomenon.

There is a limit to the length of the encoder and decoder of Survival Seq2Seq. Both networks are RNN-based. Consequently, they cannot pass information through time if their length exceeds a certain limit. However, using the attention mechanism solves this problem to a certain extent. We kept the maximum length of the encoder up to 60 time steps for MIMIC-IV. During data pre-processing, if two covariates of a data sample are measured at relatively close timestamps, we consider the average of those timestamps as the unique timestamp for those two measurements. This helps us to reduce the number of longitudinal measurements for a given row of data (or a patient) in the dataset. If the number of measurements exceeds the maximum of time steps, we only keep 60 randomly-selected measurements from that row of data, while discarding the rest. Although, we were able to use a higher maximum length, we did not notice a meaningful accuracy improvement using an encoder longer than 60 steps on MIMIC-IV. Similarly, the length of the decoder is limited as well. However, we do not believe that the limit on the length of the decoder

causes an issue for practical use cases. For example, in our experiments, we successfully trained Survival Seq2Seq on MIMIC-IV with a 2-hour time resolution. The high resolution decoder had a length twice the length of the decoder used for reporting the results in this paper. Such an experiment shows that the length of the RNN-based networks does not limit the ability of Survival Seq2Seq for practical use cases. However, one must be careful not to cause vanishing gradients by setting very lengthy encoder and decoder networks.

7. Acknowledgement

We would like to sincerely thank Health Canada for their kind support and funding the challenge "Machine learning to improve organ donation rates and make better matches" (Challenge ID: 201906-F0022-C00008). This challenge aims to improve the quality of organ matchmaking and increase the pool of Donation after Circulatory Death (DCD) donors. This study was performed to address the survival analysis aspect of the challenge.

References

- Laura Antolini, Patrizia Boracchi, and Elia Biganzoli. A time-dependent discrimination index for survival data. *Statistics in medicine*, 24(24):3927–3944, 2005.
- Yuval Barak-Corren, Victor M Castro, Solomon Javitt, Alison G Hoffnagle, Yael Dai, Roy H Perlis, Matthew K Nock, Jordan W Smoller, and Ben Y Reis. Predicting suicidal behavior from longitudinal electronic health records. *American journal of psychiatry*, 174(2):154–162, 2017.
- Viv Bewick, Liz Cheek, and Jonathan Ball. Statistics review 12: survival analysis. *Critical care*, 8(5):1–6, 2004.
- Zhengping Che, Sanjay Purushotham, Kyunghyun Cho, David Sontag, and Yan Liu. Recurrent neural networks for multivariate time series with missing values. *Scientific reports*, 8(1):1–12, 2018.
- Taane G Clark, Michael J Bradburn, Sharon B Love, and Douglas G Altman. Survival analysis part i: basic concepts and first analyses. *British journal of cancer*, 89(2):232–238, 2003.
- David Roxbee Cox and David Oakes. *Analysis of survival data*. Chapman and Hall/CRC, 2018.
- Friedo W Dekker, Renée De Mutsert, Paul C Van Dijk, Carmine Zoccali, and Kitty J Jager. Survival analysis: time-dependent effects and time-varying risk factors. *Kidney international*, 74(8):994–997, 2008.
- Matthew Engelhard, Samuel Berchuck, Joshua D’Arcy, and Ricardo Henao. Neural conditional event time models. In *Machine Learning for Healthcare Conference*, pages 223–244. PMLR, 2020.

- Shi Hu, Egill Fridgeirsson, Guido van Wingen, and Max Welling. Transformer-based deep survival analysis. In *Proceedings of AAAI Spring Symposium on Survival Prediction - Algorithms, Challenges, and Applications 2021*, volume 146 of *Proceedings of Machine Learning Research*, pages 132–148. PMLR, 22–24 Mar 2021.
- Hemant Ishwaran, Udaya B Kogalur, Eugene H Blackstone, and Michael S Lauer. Random survival forests. *The annals of applied statistics*, 2(3):841–860, 2008.
- Bingzhong Jing, Tao Zhang, Zixian Wang, Ying Jin, Kuiyuan Liu, Wenze Qiu, Liangru Ke, Y. Sun, Caisheng He, Dan Hou, Linqun Tang, xing lv, and Li Chaofeng. A deep survival analysis method based on ranking. *Artificial intelligence in medicine*, 98:1–9, 07 2019.
- A Johnson, L Bulgarelli, T Pollard, S Horng, LA Celi, and R Mark. Mimic-iv (version 1.0), 2020.
- Fahad Kamran and Jenna Wiens. Estimating calibrated individualized survival curves with deep learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 240–248, 2021.
- Jared L Katzman, Uri Shaham, Alexander Cloninger, Jonathan Bates, Tingting Jiang, and Yuval Kluger. Deepsurv: personalized treatment recommender system using a cox proportional hazards deep neural network. *BMC medical research methodology*, 18(1): 1–12, 2018.
- Faisal M Khan and Valentina Bayer Zubek. Support vector regression for censored data (svrc): a novel tool for survival analysis. In *2008 Eighth IEEE International Conference on Data Mining*, pages 863–868. IEEE, 2008.
- Changhee Lee, William R Zame, Jinsung Yoon, and Mihaela van der Schaar. Deephit: A deep learning approach to survival analysis with competing risks. In *Thirty-second AAAI conference on artificial intelligence*, 2018.
- Changhee Lee, Jinsung Yoon, and Mihaela Van Der Schaar. Dynamic-deephit: A deep learning approach for dynamic survival analysis with competing risks based on longitudinal data. *IEEE Transactions on Biomedical Engineering*, 67(1):122–133, 2019.
- Mei-Ling Ting Lee and G. A. Whitmore. Threshold Regression for Survival Analysis: Modeling Event Times by a Stochastic Process Reaching a Boundary. *Statistical Science*, 21(4):501 – 513, 2006.
- Yan Li, Jie Wang, Jieping Ye, and Chandan K Reddy. A multi-task learning formulation for survival analysis. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1715–1724, 2016.
- Chirag Nagpal, Xinyu Li, and Artur Dubrawski. Deep survival machines: Fully parametric survival regression and representation learning for censored data with competing risks. *IEEE Journal of Biomedical and Health Informatics*, 25(8):3163–3175, 2021a.
- Chirag Nagpal, Steve Yadlowsky, Negar Rostamzadeh, and Katherine Heller. Deep cox mixtures for survival regression. *arXiv preprint arXiv:2101.06536*, 2021b.

- Alfredo Nazabal, Pablo M Olmos, Zoubin Ghahramani, and Isabel Valera. Handling incomplete heterogeneous data using vaes. *Pattern Recognition*, 107:107501, 2020.
- Romain Pirracchio, Maya L Petersen, Marco Carone, Matthieu Resche Rigon, Sylvie Chevret, and Mark J van der Laan. Mortality prediction in intensive care units with the super icu learner algorithm (sicula): a population-based study. *The Lancet Respiratory Medicine*, 3(1):42–52, 2015.
- M Mostafizur Rahman and Darryl N Davis. Addressing the class imbalance problem in medical datasets. *International Journal of Machine Learning and Computing*, 3(2):224, 2013.
- Kan Ren, Jiarui Qin, Lei Zheng, Zhengyu Yang, Weinan Zhang, Lin Qiu, and Yong Yu. Deep recurrent survival analysis. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 4798–4805, 2019.
- Mohammad R Rezaei, Alex E Hadjinicolaou, Sydney S Cash, Uri T Eden, and Ali Yousefi. Direct discriminative decoder models for analysis of high-dimensional dynamical neural data. *Neural Computation*, 34(5):1100–1135, 2022.
- Donna Tjandra, Yifei He, and Jenna Wiens. A hierarchical approach to multi-event survival analysis. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 591–599, 2021.
- Bhanukiran Vinzamuri and Chandan K Reddy. Cox regression with correlation based regularization for electronic health records. In *2013 IEEE 13th International Conference on Data Mining*, pages 757–766. IEEE, 2013.
- Bhanukiran Vinzamuri, Yan Li, and Chandan K Reddy. Active learning based survival regression for censored data. In *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management*, pages 241–250, 2014.
- Zifeng Wang and Jimeng Sun. Survtrace: Transformers for survival analysis with competing events, 2021.
- Brian J Wells, Kevin M Chagin, Amy S Nowacki, and Michael W Kattan. Strategies for handling missing data in electronic health record derived data. *Egems*, 1(3), 2013.
- Kyongsik Yun, Jihoon Oh, Tae Ho Hong, and Eun Young Kim. Prediction of mortality in surgical intensive care unit patients using machine learning algorithms. *Frontiers in Medicine*, 8:406, 2021.

8. Appendix

Databases

MIMIC-IV DATABASE

The MIMIC-IV database contains health related data of ICU patients of Beth Israel Deaconess Medical Center, between 2008 and 2019. There is a total of number of 71791 distinct

Table	Total # of Covariates	# of Selected Covariates
INPUTEVENTS	282	30
OUTPUTEVENTS	69	5
CHARTEVENT	1566	73
Total	1917	108

Table 6: MIMIC-IV tables with their corresponding number of total and selected covariates.

ICU admission records with an average ICU stay of 4 days. This dataset includes vital sign measurements, laboratory test results, medications, imaging report of the patients, stored in separate tables. For mortality prediction, relevant covariates has been extracted from the following three tables: 1) INPUTEVENTS (continuous infusions or intermittent administrations), 2) OUTPUTEVENTS (patient outputs including urine, drainage, and so on), and 3) CHARTEVENTS (Patient’s routine vital signs and any additional information relevant to their care during ICU stay). We selected a total of 108 covariates from these three tables. These covariates were selected based on the feedback from our medical team, as well as conducting conventional feature selection techniques on the dataset. The number of patients after feature selection dropped to 66363 with an uncensored (deceased patients) rate of about 12%. The following table lists the total number of covariates and selected number of covariates from each table.

Data Pre-Processing Considerations

- Patients with the following diagnosis are excluded from the data: Sudden infant death synd, Unattended death, Maternal death affecting fetus or newborn, Fetal death from asphyxia or anoxia during labor, Intrauterine death.
- Invalid measurements were removed from the data using the provided WARNING and ERROR columns.
- In-valid measurements were removed from the data using the provided WARNING and ERROR columns.
- Survival time was defined as the period between the time of admission and time of death (for patients who died in hospital), or time of discharge (for censored patients).
- Time of observation is defined as the time of recording of the measurement with admission time as the baseline.