

Deep survival analysis with longitudinal X-rays for COVID-19

Michelle Shu¹ Richard Strong Bowen¹ Charles Herrmann¹ Gengmo Qi¹

Michele Santacatterina³ Ramin Zabih^{1,2}

¹ Cornell Tech ² Google Research ³ George Washington University

{michelleshu, rsb, cih, rdz}@cs.cornell.edu, gq35@cornell.edu, msanta@email.gwu.edu

Abstract

Time-to-event analysis is an important statistical tool for allocating clinical resources such as ICU beds. However, classical techniques like the Cox model cannot directly incorporate images due to their high dimensionality. We propose a deep learning approach that naturally incorporates multiple, time-dependent imaging studies as well as non-imaging data into time-to-event analysis. Our techniques are benchmarked on a clinical dataset of 1,894 COVID-19 patients, and show that image sequences significantly improve predictions. For example, classical time-to-event methods produce a concordance error of around 30-40% for predicting hospital admission, while our error is 25% without images and 20% with multiple X-rays included. Ablation studies suggest that our models are not learning spurious features such as scanner artifacts and that models which use multiple images tend to perform better than those which only use one. While our focus and evaluation is on COVID-19, the methods we develop are broadly applicable.

1. Introduction

COVID-19 is the most significant global health emergency in recent memory, with hundreds of thousands dead and widespread economic disruption. There is growing evidence that imaging is useful for the diagnosis and management of COVID-19 [1, 2]. Clinicians use radiology imaging to assess structural information which cannot be assessed with laboratory tests or physical examination. In COVID-19, chest imaging adds a high-dimensional assessment of the degree of pulmonary involvement of the disease. It allows clinicians to rule out other conditions which might contribute to the patient’s presentation such as lobular pneumonia and pneumothorax and to assess the patient for comorbidities such as heart failure, emphysema, and coronary artery disease. Some researchers have already found that imaging features predict mortality in COVID-19 [3].

In this paper we address the challenge of predicting the time course of COVID-19 patient outcomes; for example,

the probability that a specific patient will need an ICU bed in the next few days following hospital admission. Classical statistical techniques for time-to-event analysis (sometimes referred to as survival analysis) are widely used, but struggle with incorporating images due to their high dimensionality.

We begin with an overview of time-to-event analysis and a discussion of the challenges that images and COVID-19 present. Our deep learning approach is presented in section 3, followed by a review of related work in section 4. We describe our clinical dataset and some implementation details, including the baseline, in section 5. Experimental results are given in section 6, with additional data and analysis in the supplemental material.

2. Time-to-event analysis

Time-to-event analysis techniques [4] predict the probability of an outcome event occurring before a specific time, while accounting for right-censored (incompletely observed) data. Right-censoring happens when the event under study may not be observed within the relevant time period. In the clinical setting, these methods can predict a patient’s probability of undergoing an event in a particular time interval as a function of their features. In our dataset for instance, when predicting if a hospitalized COVID-19 patient will be admitted to the ICU, right-censoring happens when, as of today, the patient has not been admitted.

Time-to-event analysis focuses on three interrelated quantities: (1) the hazard function, the rate of an event at time t given that the event did not occur before time t , which is not affected by right-censoring [5]; (2) the cumulative hazard function, the integral of the hazard function between 0 and time t ; and (3) the survival function, a decreasing function that provides the probability that a patient will not experience the event of interest beyond any specified time t and is expressed as the exponential of the negative of the cumulative hazard function.

While the hazard function is not a probability, $h(t)$ can be viewed as the probability of the event occurring in a small interval around t given that the event did not occur before t . For clinical purposes, once we have estimated the hazard

function we can then compute the probability of an event occurring during a specific time interval, e.g. ICU admission in the 72 hours after hospitalization.

2.1. Cox proportional hazards model

We model the hazard function using the most popular model, the Cox model [6], defined as

$$h(t|x) = h_0(t) \exp[r(x(t))]. \quad (1)$$

Here t is the time, $x(t)$ is the set of features, $h_0(t)$ is the baseline hazard, the hazard of the specific event under study shared by all patients at time t , and $r(x(t))$ is the *risk function*, which describes the relationship between a patient's features $x(t)$ and the hazard of experiencing an event. Note that $h_0(t)$ only depends on time and not on features.

The Cox model has several advantages. First, it has no distributional assumptions, so its results will closely approximate the results for the correct model [7]. Second, even though the baseline hazard is assumed to be “unknown” and left unspecified, under minimal assumptions, the hazard, the cumulative hazard and the survival functions can be directly determined. These can then be used for predicting the probability of an event occurring before the observing time t ¹ [7].

The Cox model assumes that *time-fixed* features (features that do not change over time) have a linear multiplicative effect on the hazard function and that the hazard ratio is constant over time. This is known as the proportional hazards (PH) assumption [6]. In our task, this means that, for example, patients with a low income have a higher (or lower) hazard of dying compared with patients with a high income and this ratio is constant over time. Note that this assumption is only needed for time-fixed features and successful strategies can be easily implemented to detect and overcome its violation. Examples include the graphical approach of log-log plots for detection, as well as adding into the Cox model an interaction between the non-proportional time-fixed feature and time for overcoming its violation [9, 10].

2.2. Images present challenges

The Cox model is a mainstay of time-to-event analysis, and has been extended to deal with complex scenarios [11, 12, 13, 14, 15, 16]. However, there are two features of our task that require us to go beyond the state of the art. First, images pose a significant challenge due to their high dimensionality. Second, the time course of COVID-19 involves multiple interrelated events that cannot be predicted independently.

While there is compelling evidence that imaging studies are helpful in the diagnosis and management of COVID-19 [1, 2, 3], images present significant challenges. The amount of

¹Another statistical technique used to predict probabilities of binary events is logistic regression [8]. Although widely used, logistic regression ignores the information about the time to an event and the censoring time, while time-to-event techniques fully exploit this important information.

data in a single imaging study is orders of magnitude larger than the data available from other sources; a single medical image can easily be hundreds of megabytes.² However, the Cox proportional hazards model cannot directly handle images as features due to their high dimensionality. As [17] reports, such inputs lead to degenerate behavior. It would of course be possible to learn a feature from an imaging study, for example a rating of disease severity on a 3-point scale. Such an approach would severely and unnecessarily limit what can be learned from the images, which is a particularly poor choice for a novel disease.

As mentioned, the COVID-19 disease process involves competing and interrelated events. A straightforward application of time-to-event analysis would predict these events independently. This could easily lead to incoherent and self-contradictory predictions (for example, predicting that ICU discharge will almost certainly happen before admission to the ICU).

3. Our approach

Our main goal is to predict the probability of experiencing death, ICU admission, ICU discharge, hospital admission, and hospital discharge before the observing time t , as a function of patient's features. To do so, we assume nonlinear proportional hazards [13] for the time-fixed features.³ This assumption relaxes the more strict assumption of linear proportional hazards of the classical Cox model. In our analyses this means that, for instance, the hazard of dying among older patients at baseline increases non-linearly compared with younger patients at baseline and this ratio is constant over time [7]. This assumption has been already used in a variety of state of the art deep learning time-to-event techniques [11, 13, 18]. We also made the common assumption of non-informative censoring [19] which states that after conditioning on observable features, censoring and the risk for an event of interest are independent, *i.e.*, the censoring mechanism does not depend on the unobserved features.

To estimate the hazard function introduced in Eq. (1) we compute two components, the baseline hazard $h_0(t)$ that only depends on the time t and the risk function $r(x(t))$ that only depends on features $x(t)$. Once the hazard function is estimated, the cumulative hazard function and the survival function can be easily derived [4]; these are then used to predict the patient's probability of undergoing an event before time t . In other words, for each patient's set of features $x(t)$, we can predict the probability of an event happening before the observed time t . The baseline hazard does not depend on the features $x(t)$ and therefore, it can be easily computed by

²While CT is becoming increasingly available, early in the pandemic imaging was primarily chest X-ray (CXR).

³This assumption is only needed for time-fixed features. Time-dependent features already depend on time, making the hazard also depending on time.

using classical estimators. We used the one presented in Eq. 4.34 of [19]. The risk function, however, depends on time-fixed and time-dependent image and non-image features. To estimate the risk function $r(x(t))$ while taking into account these challenging types of features, we developed novel deep learning techniques.

3.1. Architecture

To incorporate time-dependent imaging studies, time-dependent non-image data, and time-fixed variables, our proposed architecture has three components. First, we use a convolutional LSTM (ConvLSTM) [20] and an RNN-LSTM to extract time-dependent image features and time-dependent non-image features respectively. Then, we concatenate (\oplus) the features extracted from the networks with time-fixed variables mapped to its corresponding embedding space, and passed the concatenated vector through a set of fully connected layers (FC Layers) to predict the risk function (Risk). The architecture is shown in figure 1.

3.2. Loss function

Computed on the hazard function, the Cox’s partial likelihood loss has been successfully used in recent state-of-the-art deep learning techniques [13, 18]. This likelihood function, however, is unsuitable, since it only applies to continuous time data where no two events occur at the same time. In the case of discrete time data, ties may occur and all possible orders of these tied events should be considered. Therefore, we adopt Efron’s approximation for handling ties, which is a computationally efficient estimation on the original Cox’s partial likelihood when ties are present [21].

Specifically, the loss function is as follows:

$$L = \sum_i \left[R_i - \sum_{w=0}^{d_i-1} \log \left(\sum_{j: T_j \geq t_i} e^{r(x_j(t))} - \frac{w}{d_i} U_i \right) \right],$$

where $R_i = \sum_{j \in K_i} r(x_j(t)), U_i = \sum_{j \in K_i} e^{r(x_j(t))}$. (2)

Here i denotes a unique time point and T_j is the followup time for patient j . We define the event indicator C_j , where $C_j = 1$ if the patient j experiences an event at followup time T_j , and $C_j = 0$ if censored. The risk estimate $r(x_j(t))$ is the output of our architectures for patient j . K_i is the set of patients whose followup time $T_j = t_i$ and $C_j = 1$, and $d_i = |K_i|$. It is worth noticing that by construction, this loss function does not contain the baseline hazard $h_0(t)$, making its computation easier.

3.3. Evaluation and inference

We use concordance error [22] to compare performances of different models. This calculates the proportion of falsely ordered event pairs divided by the total number of possible evaluation pairs.

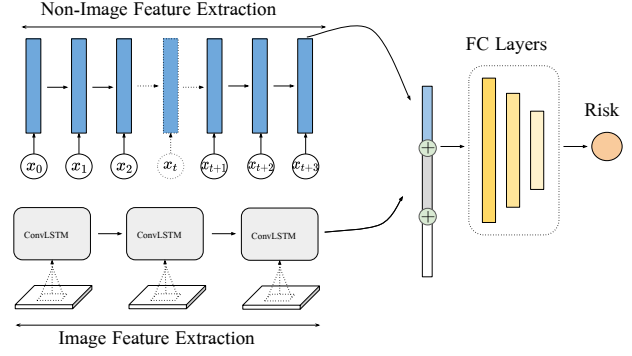


Figure 1: Our proposed architecture that handles time-fixed data, longitudinal non-image data and longitudinal images.

Similar to the loss function, during inference time, our model estimates the baseline hazard $\Delta H_0(t_i)$ and the cumulative baseline hazard $H_0(t)$ function using the Efron estimator for handling ties:

$$\Delta H_0(t_i) = \sum_{l=0}^{D_i-1} \frac{1}{\sum_{j: T_j \geq t_i} e^{r(x_j)} - \frac{l}{D_i} \sum_{j \in K_i} e^{r(x_j)}}, \quad (3)$$

$$H_0(t) = \sum_{t_i \leq t} \Delta H_0(t_i) \quad (4)$$

3.4. Mini-batch SGD and stratified sampling

In the Cox partial likelihood function, we see that the formulation involves the risk predictions of all the patients whose followup time is longer than T_i . It is computationally costly and almost infeasible to optimize this loss when models use time-dependent images. Instead, we use mini-batch stochastic gradient descent: for each iteration, we sample a subset of patients and compute the loss on the subset. To closely mimic the data distribution of the original patient group and to keep the loss function in a stable range (the range of loss function correlates to the number of events in the patient group), we use stratified sampling to maintain the same ratio of non-censored and censored patients in each mini-batch.

4. Related work

Time-to-event analysis is an important tool for many applications and plays a crucial role in healthcare. A wide variety of techniques have been developed [4], notably including non-parametric methods such as Kaplan-Meier, parametric methods such as Weibull and Gompertz, and semi-parametric models, such as the Cox model. Powerful machine learning algorithms have also been adapted to this task; popular examples include Survival Support Vector Machines [23] and Random Survival Forests [24]. The Cox model, which directly estimates the hazard function, stands out as perhaps the most popular time-to-event analysis method.

Deep Survival Analysis. In recent years, deep learning has been used to extend the Cox model. DeepSurv [13] was one of the first to use deep learning to model the risk function $r(x(t))$ in the hazard function. They demonstrated strong performance in both linear and non-linear settings. There are also attempts to tackle the limitations of DeepSurv to only structured data. By using convolution neural networks, [18, 25, 26] models hazards on unstructured features such as images, which are much harder to incorporate due to their high dimensionality. Our architecture further extends the Cox model so that it takes both structured and unstructured longitudinal input.

Unstructured Longitudinal Data. Incorporating longitudinal data in time-to-event prediction has received increasing attention, typically using deep learning techniques. [14] proposes DeepHit, a model that enables the possibility that the relationship between covariates and risks to change over time. [27, 28, 29] have demonstrated the effectiveness of recurrent neural networks on longitudinal structured datasets. We appear to be the first to adopt recurrent neural network on longitudinal unstructured medical data, and the recurrent nature of our model is successful in capturing time dependency relationships within the data.

End-to-End Training. As mentioned in section 3.4, the negative partial likelihood function loss function is computationally expensive and almost infeasible (in terms of GPU memory) to compute when we have a large patient group with unstructured image data. [30] adopted an alternative two-step training strategy which requires first training a feature extraction network on image data with expert labelling, then uses the extracted features as covariates of the Cox model. Using mini-batch sampling, our training process is end-to-end and does not require any expert labelling.

5. Experimental setup and clinical dataset

Our techniques are designed for a clinical setting, where a combination of time-dependent and time-independent patient features are available, including imaging. There are no publicly available datasets for COVID-19 that contain this information, and patient privacy considerations make it unlikely that such data will be available anytime soon. The closest existing datasets for COVID-19 focus primarily on images, and generally do not contain significant additional information. The recent BIMCV-COVID19+ dataset [31] is an exception, and contains a limited amount of information such as demographics and antibody test results, but falls far short of the detailed clinical information that our methods are designed to exploit. Notably, it does not contain patient outcomes or lab values.

5.1. Clinical dataset

We are fortunate to have IRB-approved access to a large clinical dataset of 1,894 COVID-19 patients, longitudinally

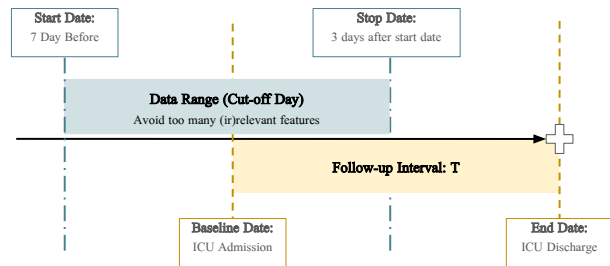


Figure 2: A sample timeline for ICU Discharge event prediction. The baseline date for ICU Discharge is ICU Admission. To avoid giving models too much information, we set the data range of this event to be 10 days. Any data before and after this range will be removed.

collected from 90 consecutive days early in the pandemic. The dataset includes patient demographics, hourly-recorded vital signs, treatment regimes and clinical notes. Crucially for our purposes there are a significant number of imaging studies also available: over 14,000 studies, primarily chest X-ray (CXR) but also an increasing number of computed tomography (CT) exams. Outcomes are also available, including such important events as hospitalization admission and discharge, ICU admission and discharge, and mortality. All patients tested COVID-19 positive by PCR. Multiple X-rays, 4 types of time-fixed features and 29 types of time-dependent non-image features of patients during their stay are also available. Time-fixed features are smoking and pregnancy status, active cancer and age at admission.

The data was provided to us by a study research coordinator. Due to possible inconsistencies in the ways that events are recorded at their institution, there may be a small number of patients whose date of events was not included. In such a case, the patient's events occurred but were unobserved, which are right-censored data in a time-to-event framework. Our approach to censored outcomes, discussed below, handles this situation straightforwardly.

We process the available data with the following pipeline:

- For each type of event, we select a baseline date for the Cox model based on advice from clinicians, and generate a time interval for each patient. The start date and end (followup) date of the interval is dependent on the baseline date and the type of events. We filter out any patient without an appropriate baseline date.
- We discretize the time interval generated for each qualified patient by day, and associate each time-dependent feature in this time interval with their corresponding days. We remove patients who do not have any X-rays taken during this interval.

	Admit	ICU Admit	ICU Discharge	Discharge	Death
Event #	137	287	251	1171	290
Censored #	99	224	150	219	427
Event Removed #	1395	138	1	26	6
Censored Removed #	263	1245	1492	478	1171
Baseline Date	First available X-ray	Admit	ICU Admit	Admit	Admit
Start Date	Same as Baseline	7 day Before	7 Day Before	7 Day Before	7 Day Before
Cut-Off Day	30 Days	10 Days	10 Days	10 Days	10 Days

Table 1: Details on the data distribution for each type of events.

	Age	Smoking	Pregnancy	Cancer
Data Breakdown	IQR: 23 Median: 64	Active Smoker: 70 Former Smoker: 382 Non-Smoker: 1257	Yes: 7 No: 1705	Solid: 71 Liquid: 42 No: 1257
Missing Report	29.8%	9.8%	9.6%	27.6%

Table 2: Time-fixed features in our clinical dataset.

- A cut-off day is selected for each type of event, and we remove any data collected beyond this cut-off day.

Table 1 summarizes our choices of baseline date, start date and cut-off day for each type of event. For subjects with event outcomes, the date of the followup is exactly the date when the event occurred. For censored outcomes (including unrecorded deaths), we use the date of the most recent X-ray to be the followup date, where the study terminates without reaching a definite conclusion. Statistics on time-fixed features are provided in Table 2. For completeness, the set of time-dependent features can be found in the supplements.

In our dataset, longitudinal information is available until the patient reaches an absorbing state. This is not a realistic setting, since predictions are mostly likely needed at earlier stages, where only data close to the baseline date are available. However, our cut-off day procedure prevents simple cheating such as counting the number of days where data is available. The cut-off also avoids other forms of cheating, for example if we go far enough beyond the baseline date, some lab results may start to give obvious clues.

Note that our model is recurrent and capable of analyzing input without any X-rays. However, we want to study the impact of images on time-to-event analysis, and wish to compare our model with another image-based survival analysis method – DeepConvSurv [18]. As a result we have filtered out patients who did not have any X-rays taken during the corresponding intervals.

After obtaining the desired patient population for each event, we follow the standard train, validation, test split (60:20:20). We divide the population to train, validation and test set such that these subsets share the same event and censored ratio.

5.2. Implementation details

Non-image time-dependent data is represented by a matrix of size $D \times 2W$, where D is the cut-off day, i.e. the maximum number of days of data the model could use in this event prediction, and W is the total number of types of lab results available. For each vector of length $2W$ in this matrix, the first W values x_1, \dots, x_W are the lab results re-scaled to 0-1 range, and the remaining W values y_1, \dots, y_W are indicators. $y_i = 1$ if the corresponding x_i is a lab measurement available on that day, and 0 otherwise.

X-ray images of the same patient are downsized and stacked, resulting in a tensor of size $224 \times 224 \times O$, where O is the number of patient X-rays available for this event.

To predict the risk function $r(x_j)$ for patient j , the image data is fed to a convLSTM (kernel size 7×7 , 2 hidden layers size 32 and 16) recurrently to extract image features. The extracted features are further downsized by a 2D Adaptive Average Pooling layer (2×2) following a fully connected layer (size 64). Non-image time-dependent inputs are first mapped to an embedding space (embedding size 15), then fed recurrently to LSTM to extract the respective features. Patient demographics, which are all discrete categorical values, are first mapped to their corresponding embedding space (embedding size 2), then concatenated altogether with the extracted features from convLSTM and LSTM branches. The concatenated vector goes through three fully connected layers (size 32, 16 and 1) to predict the final outcome.

During training, we use Adam [32] with $lr = 1e - 3$, with a batch size of 40. We divide the training phase into 30 epochs, each epoch consists of 20 randomly sampled mini-batches. With limited GPU memory available and to reduce computational costs, we further sample the X-ray data and limit the number of X-rays k to $k = 4$ per patient, and use

zero padding when patients have less than k valid X-rays.

In the validation phase, for each epoch we choose the mini-batch that the model has the lowest concordance error on as our baseline hazard, and compute the concordance error on the validation set. The epoch with the lowest concordance error and its best-performing mini-batch on validation set will be used to compete against other baseline models on the testing set.

5.3. Baseline details

We compare our proposed architecture with 8 different models, including parametric and semi-parametric models, standard and non-linear CoxPH models, and other popular machine learning survival methods:

- **Non-Image Input Models:** Weibull, Gompertz [33], Survival Forest [34], Survival SVM [35], the standard Cox Proportional Hazard model (CoxPH) [6], the non-linear CoxPH model (DeepSurv) [13]
- **Image Input Models:** Deep Convolutional Neural Network for Survival Analysis (DeepConvSurv) [18]

Recall that for our model, the non-image time-dependent input are formatted to matrix of size $D \times 2W$. For the non-image baselines, which take 1D input, we concatenate the row vectors in the matrix along the day axis, and further combine it with time-fixed data to obtain a long vector representation of the patient data.

For DeepSurv [13], we used a model with 2 fully connected layers (size 128, 64) with ReLu activation. For DeepConvSurv [18], we implemented a network very similar to the original, which consists of three convolution layers (conv1: kernel size 7×7 , stride 3, channel 32; conv2: kernel size 5×5 , stride 2, channel 32; conv3: kernel size 3×3 , stride 2, channel 32), following a single FC layer. We also added max pooling and 2D dropout to the network to improve model performance. Similar to our own architecture, we train DeepConvSurv using randomly selected images for each patient, and experiment on both the case where the loss is computed over the entire training set, and the case where mini-batch gradient descent is used.

To investigate the impact of images on prediction, we provide two baselines of our own: 1) a model where the convLSTM branch is removed and predicts only from non-image data, and 2) an “complementary” model where the LSTM branch is removed and we predict solely from image-data. In these two baseline models, we also use mini-batch gradient descent with batch size of 40.

6. Results

6.1. Performance comparisons

We evaluate model performance using time-dependent concordance error discussed in section 3.3. Results of the




Figure 3: Brier score comparison against selected standard techniques, lower is better

comparisons are in Table 3. Additional experiments and ablation studies are included in the supplemental material.

In our experiments, non-linear CoxPH models (DeepSurv and ours) and Random Survival Forest almost always outperform parametric models. This is expected since Cox models make no distributional assumptions about the data. The performances of DeepConvSurv and our Image-Only baseline are also noteworthy: they demonstrate that image alone (without any lab values or demographics) provides useful information for prediction. Using multiple images further improves the performances.

Our recurrent baseline model, which captures time-dependency relations, achieves competitive performance across all 5 events even without images. Adding our ConvLSTM branch that processes time-dependent images has further improved the predictions by an average of 2.8%, with a substantial improvement of 4.9% on Hospitalization (Admit) and 5.4% on ICU Discharge.

To better illustrate the comparison, in figure 3 we show the time-dependent Brier score [36] for our method and two standard time-to-event prediction techniques. This is an extension of the Brier score [37] for a specific time horizon.

Our experiments demonstrate the effectiveness of our recurrent architecture on non-linear, longitudinal data. We also show that incorporating multiple, time-dependent imaging studies significantly improves time-to-event predictions.

6.2. Understanding the dataset and model

We conduct several experiments to provide insights into what our models learned.

Feature Importance Test: While our imaging studies were anonymized, fields describing individual scanners were preserved. This provides some ways to check if our model is simply learning an association between certain scanners and disease severity (for example, sicker patients might be in certain areas of the hospital, and routed to the nearest x-ray machine). To test whether our model is learning

Method	Admit	ICU Admit	ICU Discharge	Discharge	Death
CoxPH	0.434	0.371	0.486	0.240	0.358
Weibull	0.319	0.372	0.478	0.437	0.477
Gempertz	0.382	0.362	0.413	0.272	0.360
DeepSurv	0.316	0.358	0.454	0.398	0.373
Survival+SVM	0.333	0.365	0.464	0.347	0.316
Survival+Random Forest	0.337	0.309	0.435	0.427	0.282
DeepConvSurv (GD)	0.403	0.417	0.468	0.356	0.351
DeepConvSurv (mini-batch)	0.419	0.415	0.478	0.419	0.435
Image-Only (ours)	0.238	0.408	0.459	0.350	0.401
Non-Image (ours)	0.247	0.265	0.439	0.262	0.278
Image+Non-Image (ours)	0.198	0.241	0.385	0.229	0.246

Table 3: Concordance error for time-to-event predictions. Lower is better, colors encode the best performing 3 methods for each event. DeepConvSurv with Gradient Descent (GD) computes its loss function over the entire patient group, while DeepConvSurv with mini-batch Stochastic Gradient Descent (mini-batch) adopts the same sampling strategy as our model described in section 3.4.

simple associations between scanners and events, we conducted permutation-based feature importance tests (using survival random forest) to measure performance drops if certain features are removed. Specifically, we add scanner IDs of patient x-rays as one of the features for event prediction, and measure the average concordance error increase when the relationship between the feature and survival time is removed by random shuffling. We found that for Hospital Admission and ICU Admission, Scanner ID has an average concordance error increase of 0.0020 ± 0.0044 and 0.0060 ± 0.0703 respectively, and 0.00 ± 0.000 on the ICU Discharge, Discharge and Death events. Compared with other features such as Age (0.0682 ± 0.0542 on ICU Discharge), the permutation tests results on Scanner ID suggest little correlation between scanner types and events.

Black Box Test: Recent papers, such as [38], have identified flaws in the evaluation of deep learning models for COVID-19. Deep learning models were able to utilize non-medical information in the X-rays like artifacts or watermarks in order to differentiate between patients with COVID and those without. In order to demonstrate that our dataset and model do not exhibit this flaw, we performed similar experiments to [38], where we obscured various portions of the images with black boxes and retrained the model. This should determine what part of the images are important to the model. If our technique is relying on incidental properties of the images, such as identifiable artifacts from individual scanners, we would expect to see better prediction accuracy from an image and labs model over a labs-only model even when most of the image is obscured.


Results are shown in table 4. The first and last columns give our results with images included (first column) and non-image data only (last column). We used several different sized boxes and zeroed out the portions of the image within them. As a sanity check we obscured all of the image

(“100% obscured”). We also obscured large portions of the image (“80% obscured”, “90% obscured”, the percentages indicate how much of the image was obscured starting at the center). Note that sometimes clues concerning illness severity (e.g. tubes) are located at the top, and thus visible in the 80% obscure and 90% obscure cases. Therefore, we also included a “95% cols obscured” case, where 95% of the center columns are removed. The boxes are shown overlaid on a publicly available chest x-ray above table 4.

The experimental results suggest that our models are not learning coincidental features such as scanner artifacts or watermarks. We also observe that there are some small improvements when our models use 80% obscured images, and there is a natural explanation: unless the image is almost entirely obscured, various lines and sensors are easily discernible even outside the lungs, and these possibly provide an indication of disease severity.

Number of Input Images Ablation: We also perform an ablation on the number of images used (k) when training and evaluating the model. We find that the model improves as the number of images increases for the majority of tasks (admission, ICU discharge, and discharge). Specifically, we observe that $k = 2$ provides a 10-20% improvement in concordance error over $k = 1$; for example, going from $k = 2$ to $k = 1$ leads to a decrease in concordance error of 0.072 in admission; 0.006 in discharge; and 0.048 in ICU discharge. The remaining tasks show little to no improvement from a larger number of images. However, an analysis of the dataset statistics show that these tasks tend to have fewer images per patient; this means that our training set also has fewer examples of patients with more data which may leads to problems with generalization. We include more details on this ablation in the supplemental material.

Mini-batch Size Ablation: We also investigated the effects of mini-batch size B , and observed that mini-batch



Event	Image + labs (ours)	80% obscured	90% obscured	95% cols obscured	100% obscured	Labs only (ours)
Admit Date	0.198	0.250	0.257	0.287	0.268	0.247
ICU Admit	0.241	0.262	0.271	0.262	0.264	0.265
ICU Discharge	0.335	0.413	0.437	0.431	0.437	0.439
Discharge	0.229	0.272	0.268	0.276	0.264	0.262
Death	0.254	0.296	0.274	0.271	0.278	0.278

Table 4: Ablation on obscuring parts of the input images. The images on the top provide a visual example for each input type; note that the portion inside the blue box is set to zero. (Original x-ray image taken from a publicly available dataset.) The table shows concordance error for time-to-event predictions (lower is better). If our model was learning coincidental features outside the patient’s body, such as scanner artifacts, we would expect that obscuring the center of the image would result not significantly degrade performance, which is not the case.

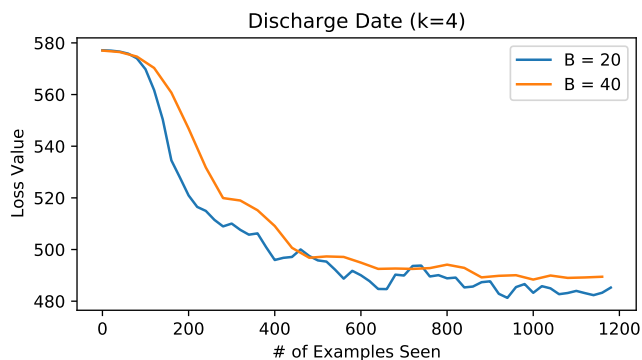


Figure 4: Validation error curves for models trained with $k = 4$ images per patients and mini-batch size $B = 20$ or 40 respectively. The convergence properties do not seem to be affected by mini-batch size.

size does not affect the model performance for sufficiently large values of B ($B \geq 20$). Particularly small mini-batches ($B = 10$) contains very few events, which results in performance drops. However, for $B \geq 20$, the models tend to converge to the same point. In figure 4, we show an example of the validation error curves of discharge events with respect to the number of examples seen for the $k = 4$ case, where $B = 20$ and $B = 40$. From the figure, we see that the size of the mini-batch does not seem to affect convergence.

7. Conclusions

We describe a deep learning approach that incorporates multiple, time-fixed data, longitudinal non-image data and

longitudinal images into time-to-event analysis. Our technique accurately predicts the probability of experiencing an event in the presence of right-censored data.

We used a large COVID-19 dataset containing longitudinal imaging and non-imaging information. While this dataset contains valuable information to predict the occurrence of clinical events, there is some risk of selection bias. For instance, we only included in our analysis patients with multiple imaging studies over time. Multiple images are usually performed when a patient is sicker, for example to confirm central line placement. This selection could lead to a sample that is not representative of the COVID-19 hospitalized population. Selection bias is a common problem in machine learning [39], statistics [40], and epidemiology [41]; as a result, a number of techniques have been developed to correct it [39].

We have demonstrated that neural networks can be used to explicitly support transitions between competing events and be used to predict the transition-specific risk (e.g., the risk of transition from hospitalization to ICU admission) for a particular set of patient features. While our focus is on COVID-19, the techniques we propose should be generally applicable to a wide class of serious diseases where imaging can improve the prediction of patient outcomes.

Acknowledgements We thank Joshua Geleris MD and George Shih MD for their help understanding various clinical issues. We received considerable help acquiring the dataset from Martin Prince MD PhD (PI on the IRB), Benjamin Cobb, Sadjad Riyahi MD, and Evan Sholle. This research was supported by a gift from Sensetime.

References

- [1] Damiano Caruso, Marta Zerunian, Michela Polici, Francesco Pucciarelli, Tiziano Polidori, Carlotta Rucci, Gisella Guido, Benedetta Bracci, Chiara de Dominicis, and Andrea Laghi. Chest CT Features of COVID-19 in Rome, Italy. *Radiology*, page 201237, April 2020. 1, 2
- [2] Sakiko Tabata, Kazuo Imai, Shuichi Kawano, Mayu Ikeda, Tatsuya Kodama, Kazuyasu Miyoshi, Hirofumi Obinata, Satoshi Mimura, Tsutomu Koderu, Manabu Kitagaki, Michiya Sato, Satoshi Suzuki, Toshimitsu Ito, Yasuhide Uwabe, and Kaku Tamura. Clinical characteristics of COVID-19 in 104 people with SARS-CoV-2 infection on the Diamond Princess cruise ship: A retrospective analysis. *The Lancet Infectious Diseases*, June 2020. 1, 2
- [3] Danielle Toussie, Nicholas Voutsinas, Mark Finkelstein, Mario A Cedillo, Sayan Manna, Samuel Z Maron, Adam Jacobi, Michael Chung, Adam Bernheim, Corey Eber, Jose Concepcion, Zahi Fayad, and Yogesh Sean Gupta. Clinical and Chest Radiography Features Determine Patient Outcomes In Young and Middle Age Adults with COVID-19. *Radiology*, page 201754, May 2020. 1, 2
- [4] David G Kleinbaum and Mitchel Klein. *Survival Analysis*, volume 3. Springer, 2010. 1, 2, 3
- [5] Jan Beyersmann, Arthur Allignol, and Martin Schumacher. *Competing Risks and Multistate Models with R*. Springer Science & Business Media, 2011. 1
- [6] David R Cox. Regression models and life-tables. *Journal of the Royal Statistical Society: Series B (Methodological)*, 34(2):187–202, 1972. 2, 6
- [7] David G Kleinbaum and Mitchel Klein. The Cox proportional hazards model and its characteristics. In *Survival Analysis*, pages 97–159. Springer, 2012. 2
- [8] David G Kleinbaum, K Dietz, M Gail, Mitchel Klein, and Mitchell Klein. *Logistic Regression*. Springer, 2002. 2
- [9] David G Kleinbaum and Mitchel Klein. Extension of the Cox proportional hazards model for time-dependent variables. In *Survival Analysis*, pages 241–288. Springer, 2012. 2
- [10] David G Kleinbaum and Mitchel Klein. Evaluating the proportional hazards assumption. In *Survival Analysis*, pages 161–200. Springer, 2012. 2
- [11] David Faraggi and Richard Simon. A neural network model for survival data. *Statistics in medicine*, 14(1):73–82, 1995. 2
- [12] Robert Tibshirani. The lasso method for variable selection in the Cox model. *Statistics in medicine*, 16(4):385–395, 1997. 2
- [13] 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):24, 2018. 2, 3, 4, 6
- [14] 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. 2, 4
- [15] Jie Hao, Youngsoon Kim, Tejaswini Mallavarapu, Jung Hun Oh, and Mingon Kang. Cox-PASNet: Pathway-based sparse deep neural network for survival analysis. In *2018 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pages 381–386. IEEE, 2018. 2
- [16] 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 bio-medical engineering*, 67(1):122–133, January 2020. 2
- [17] Noah Simon, Jerome Friedman, Trevor Hastie, and Rob Tibshirani. Regularization paths for Cox’s proportional hazards model via coordinate descent. *Journal of statistical software*, 39(5):1, 2011. 2
- [18] Xinliang Zhu, Jiawen Yao, and Junzhou Huang. Deep convolutional neural network for survival analysis with pathological images. In *2016 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pages 544–547. IEEE, 2016. 2, 3, 4, 5, 6
- [19] John D Kalbfleisch and Ross L Prentice. *The Statistical Analysis of Failure Time Data*, volume 360. John Wiley & Sons, 2011. 2, 3
- [20] Xingjian Shi, Zhoung Chen, Hao Wang, Dit-Yan Yeung, Wai-kin Wong, and Wang-chun WOO. Convolutional LSTM Network: A Machine Learning Approach for Precipitation Nowcasting. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 802–810. Curran Associates, Inc., 2015. 3

- [21] Bradley Efron. The efficiency of Cox’s likelihood function for censored data. *Journal of the American statistical Association*, 72(359):557–565, 1977. 3
- [22] Hajime Uno, Tianxi Cai, Michael J. Pencina, Ralph B. D’Agostino, and L. J. Wei. On the C-statistics for evaluating overall adequacy of risk prediction procedures with censored survival data. *Statistics in Medicine*, 30(10):1105–1117, May 2011. 3
- [23] 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. 3
- [24] Leo Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001. 3
- [25] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. Unpaired Image-To-Image Translation Using Cycle-Consistent Adversarial Networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2223–2232, 2017. 4
- [26] Alice Zheng and Amanda Casari. *Feature engineering for machine learning: principles and techniques for data scientists.* ” O’Reilly Media, Inc.”, 2018. 4
- [27] Jiarui Jin, Yuchen Fang, Weinan Zhang, Kan Ren, Guorui Zhou, Jian Xu, Yong Yu, Jun Wang, Xiaoqiang Zhu, and Kun Gai. A deep recurrent survival model for unbiased ranking. *arXiv preprint arXiv:2004.14714*, 2020. 4
- [28] Eleonora Giunchiglia, Anton Nemchenko, and Mihaela van der Schaar. Rnn-surv: A deep recurrent model for survival analysis. In *International Conference on Artificial Neural Networks*, pages 23–32. Springer, 2018. 4
- [29] 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. 4
- [30] Yifan Peng, Tiarnan D Keenan, Qingyu Chen, Elvira Agrón, Alexis Allot, Wai T Wong, Emily Y Chew, and Zhiyong Lu. Predicting risk of late age-related macular degeneration using deep learning. *NPJ digital medicine*, 3(1):1–10, 2020. 4
- [31] Maria de la Iglesia Vaya, Jose Manuel Saborit, Joaquim Angel Montell, Antonio Pertusa, Aurelia Bustos, Miguel Cazorla, Joaquin Galant, Xavier Barber, Domingo Orozco-Beltran, Francisco Garcia-Garcia, Marisa Caparros, German Gonzalez, and Jose Maria Salinas. BIMCV COVID-19+: a large annotated dataset of RX and CT images from COVID-19 patients, 2020. See: <https://bimcv.cipf.es/bimcv-projects/bimcv-covid19/>. 4
- [32] Diederick P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *ICLR*, 2015. 5
- [33] David G Kleinbaum and Mitchel Klein. Parametric survival models. In *Survival Analysis*, pages 289–361. Springer, 2012. 6
- [34] Hemant Ishwaran, Udaya B Kogalur, Eugene H Blackstone, Michael S Lauer, et al. Random survival forests. *The annals of applied statistics*, 2(3):841–860, 2008. 6
- [35] Vanya Van Belle, Kristiaan Pelekman, JAK Suykens, and Sabine Van Huffel. Support vector machines for survival analysis. In *Proceedings of the Third International Conference on Computational Intelligence in Medicine and Healthcare (CIMED2007)*, pages 1–8, 2007. 6
- [36] R Schoop, E Graf, and M Schumacher. Quantifying the predictive performance of prognostic models for censored survival data with time-dependent covariates. *Biometrics*, 64(2):603–610, 2008. 6
- [37] Erika Graf, Claudia Schmoor, Willi Sauerbrei, and Martin Schumacher. Assessment and comparison of prognostic classification schemes for survival data. *Statistics in medicine*, 18(17-18):2529–2545, 1999. 6
- [38] Gianluca Maguolo and Loris Nanni. A critic evaluation of methods for covid-19 automatic detection from x-ray images. *arXiv preprint arXiv:2004.12823*, 2020. 7
- [39] Corinna Cortes, Mehryar Mohri, Michael Riley, and Afshin Rostamizadeh. Sample selection bias correction theory. In *International conference on algorithmic learning theory*, pages 38–53. Springer, 2008. 8
- [40] Alice S Whittemore. Collapsibility of multidimensional contingency tables. *Journal of the Royal Statistical Society: Series B (Methodological)*, 40(3):328–340, 1978. 8
- [41] James M Robins. Data, design, and background knowledge in etiologic inference. *Epidemiology*, pages 313–320, 2001. 8