

# TIME-TO-EVENT PRETRAINING FOR 3D MEDICAL IMAGING

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

With the rise of medical foundation models and the growing availability of imaging data, scalable pretraining techniques offer a promising way to identify imaging biomarkers predictive of future disease risk. While current self-supervised methods for 3D medical imaging models capture local structural features like organ morphology, they fail to link pixel biomarkers with long-term health outcomes due to a *missing context problem*. Current approaches lack the temporal context necessary to identify biomarkers correlated with disease progression, as they rely on supervision derived only from images and concurrent text descriptions. To address this, we introduce *time-to-event pretraining*, a pretraining framework for 3D medical imaging models that leverages large-scale temporal supervision from paired, longitudinal electronic health records (EHRs). Using a dataset of 18,945 CT scans (4.2 million 2D images) and time-to-event distributions across thousands of EHR-derived tasks, our method improves outcome prediction, achieving an average AUROC increase of 23.7% and a 29.4% gain in Harrell’s C-index across 8 benchmark tasks. Importantly, these gains are achieved without sacrificing diagnostic classification performance. This study lays the foundation for integrating longitudinal EHR and 3D imaging data to advance clinical risk prediction.

## 1 INTRODUCTION

Foundation models for medical imaging have the potential to transform healthcare by assisting doctors in complex clinical decision making (Saab et al., 2024; Sox et al., 2024) and identifying novel pixel biomarkers predictive of future disease risk (Pai et al., 2024; Sriram et al., 2021). Such models rely on self-supervised learning (SSL) for obtaining supervision at the scale required to train on growing collections of 3D data (e.g., CT scans, MRIs). SSL captures local structural features by leveraging pretraining signal directly from images or cross-modal pairs (e.g., images and their text descriptions) (Zhang et al., 2022). While excelling at segmentation tasks and diagnostic classification of pathologies (Pinto-Coelho, 2023), SSL fails to learn prognostic biomarkers because the current pretraining regimens suffer from a *missing context problem* (see Figure 1). This issue arises when supervision sources are restricted to narrow time windows around the image, thus excluding long-term temporal patterns that are correlated with disease progression, which limits a model’s ability to identify prognostic biomarkers.

These long-term temporal patterns, which we refer to as longitudinal context, are readily available in a patient’s electronic health record (EHR) which is routinely used by clinicians to guide the interpretation of images and inform treatment planning (Leslie et al., 2000; Holste et al., 2024). Longitudinal EHRs contain temporal information about the progression of disease, as well as years of patients’ health outcomes. However, the full breadth of these outcome data, in terms of temporal structure and task diversity, is rarely used as a source of supervision when training image foundation models. Current approaches for training 3D imaging models typically restrict labels to diagnosis codes sourced from the same or nearby temporal context as the image and its textual description (Blankemeier et al., 2024). Because EHR data is readily available, it offers an untapped resource for large-scale pretraining of medical image models in a manner that uses long-term temporal context. More generally, image foundation models must reflect the settings in which they will be used, which is to assist in prognosis in the clinic (Negro-Calduch et al., 2021; Yala & Hughes, 2023).

054 However, performing effective risk estimation using imaging models involves navigating several  
 055 challenges. Developing such models requires capturing correlations between pixels and outcomes  
 056 spanning years, which is difficult with current SSL methods. Direct approaches to identifying pixel-  
 057 level biomarkers, such as sequential image capturing (Lu et al., 2021; Bera et al., 2022), are difficult  
 058 to scale for collecting large, high-quality datasets. Moreover, doing risk estimation requires address-  
 059 ing right censoring, where the outcome of interest remains unobserved by the study’s end. Naively  
 060 excluding censored patients introduces bias and reduces the available training data.

061 In medicine, time-to-event (TTE) modeling (also known as survival modeling) is commonly used  
 062 to estimate future risk of an outcome at a specific time point conditioned on feature representations.  
 063 Although TTE models offer many theoretical advantages, including the ability to estimate instanta-  
 064 neous risk at any given time point (Collett, 2023) and naturally handling right censoring (Kleinbaum  
 065 & Klein, 1996), their use in image pretraining remains underexplored. Prior deep learning studies  
 066 exploring TTE modeling in medical imaging have been restricted to small-scale, single-task appli-  
 067 cations, typically using 2D, end-to-end models (Zhu et al., 2016; Shu et al., 2021; Lu et al., 2021),  
 068 and we argue 3D medical imaging data has richer biomarkers and thus better suited for clinical de-  
 069 cision tools, especially for opportunistic detection of underdiagnosed conditions (Aali et al., 2024).  
 070 However, large-scale TTE pretraining for 3D imaging has not yet been investigated, likely because  
 071 multimodal medical datasets linking 3D images with longitudinal EHR data have only recently be-  
 072 come available (Huang et al., 2024).

073 In this work, we propose *time-to-event pretraining* for 3D medical imaging models as a way to  
 074 address the missing context problem. Our central claim is that temporal supervision, defined by  
 075 TTE distributions sourced from longitudinal EHRs, provides a readily available, scalable source of  
 076 contextual information for pretraining that better captures prognostic pixel biomarkers. Moreover,  
 077 by naturally handling right-censorship, TTE-based methods improve data efficiency and mitigate  
 078 censorship bias. Our contributions are as follows:

- 079 • We present the first large-scale evaluation of time-to-event pretraining for 3D medical imaging  
 080 encoders. We use a public dataset of 18,945 chest CT scans (equivalent to 4.2 million 2D images)  
 081 linked to longitudinal EHR data containing 225M clinical events with a median follow-up time of  
 082 5 years.
- 083 • Our approach converts longitudinal EHR data into a source of time-to-event supervision, thus  
 084 predicting not only if a clinical event will occur but also when. This richer pretraining signal goes  
 085 beyond diagnostic classification explored in prior work and enables generating many pretrain-  
 086 ing tasks (8,192 in this work) that capture the temporal event structure available in longitudinal  
 087 EHR data. This choice also increase per-image label density by an average of 3 times over prior  
 088 approaches.
- 089 • Our approach substantially improves performance in predicting future medical outcomes, achiev-  
 090 ing on average a 23.7% increase in AUROC and a 29.4% improvement in Harrell’s C-index over  
 091 baseline models for 8 benchmark tasks without negatively impacting diagnostic classification  
 092 performance in 8 external tasks. Our approach also improves model calibration, measured by the  
 093 Integrated Brier Score, by an average of 54%.

094 All our experiments are conducted using public medical datasets to ensure full reproducibility. We  
 095 also make all of our experiment code and pretrained model checkpoints available for download<sup>1,2</sup>, to  
 096 contribute to the community for continued pretraining of imaging foundation models with prognosis  
 097 as added benefit.

## 099 2 RELATED WORK

101 **Time-to-Event Modeling with Medical Images** Time-to-event (TTE) modeling, also known as  
 102 survival analysis, predicts the distribution of time until a specific event occurs, such as death. TTE  
 103 models primarily include accelerated failure time models, which assume various probability distri-  
 104 butions, e.g., exponential (Saikia & Barman, 2017), Weibull (Breheny, 2015), and Cox proportional  
 105 hazard (Cox-PH) models (Cox, 1972) which is a semi-parametric approach with a constant hazard

106 <sup>1</sup>Redacted for review

107 <sup>2</sup>Redacted for review

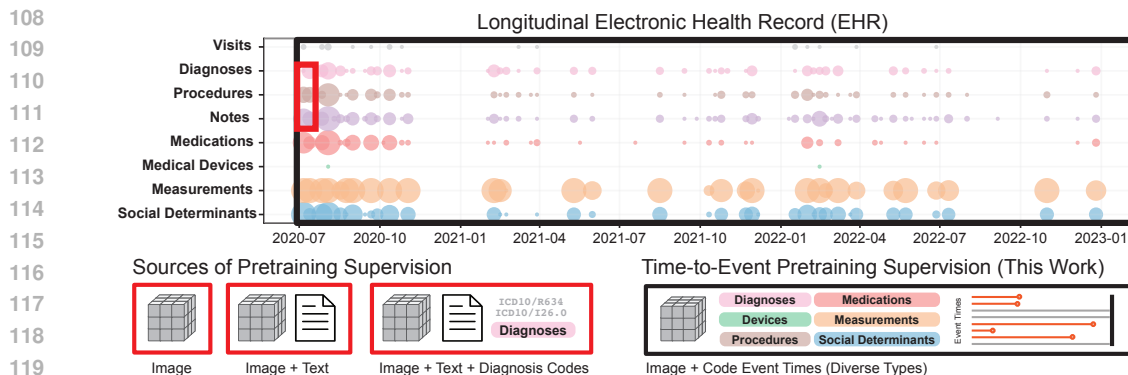


Figure 1: The *missing context problem* in medical imaging. Existing supervision sources (red boxes) are localized to the image itself (i.e., pixel features and descriptions of those features via text) or immediate clinical context via diagnosis codes. Doing so misses future information on disease progression (black boxes), which reduces the ability to learn correlations necessary for identifying prognostic pixel biomarkers. *Time-to-event pretraining* provides a principled framework for incorporating the vast amount of temporal supervision available in EHR data to estimate future risk in the presence of right censorship as well as leverage a large, diverse number of clinical tasks, beyond just diagnoses, for pre-training.

ratio assumption. Non-parametric models like random survival forests (Ishwaran et al., 2008) capture non-linear interactions. In the deep learning era, methods such as DeepSurv (Katzman et al., 2018) and MOTOR (Steinberg et al., 2024) provide higher level feature learning, making TTE modeling easier to extend to complex inputs such as medical images. Prior TTE methods for imaging assume 2D and 2.5D model architectures and have focused on end-to-end training for small-scale, single-task models. DeepConvSurv (Zhu et al., 2016) was the first work to replace the log-partial hazard (the exponential component of a Cox model) with a CNN, enabling survival prediction directly from 2D images. Similarly, Shu et al. (2021) and Lu et al. (2021) modeled the log-partial hazard with a CNN-RNN to encode sequential images.

**Pretraining for 3D Medical Image Models** Although early work in medical imaging relied on supervised pretraining using general-domain datasets such as ImageNet (Xie & Richmond, 2018; Ke et al., 2021), self-supervised learning (SSL) is now the predominate approach to scaling pretraining in medical imaging models. Popular approaches include reconstruction or de-noising objectives via masked autoencoders (MAE) (He et al., 2022), contrastive losses defined over paired samples (Chen et al., 2020; Sowrirajan et al., 2021) or leveraging multimodal pairs, such as medical images and their aligned text descriptions (Radford et al., 2021; Zhang et al., 2022) or unpaired medical images and text (Wang et al., 2022). SSL methods face added challenges when extended to 3D imaging, where instances (e.g., CT scans) contain over 100 times more pixel data than 2D counterparts like radiographs. Tang et al. (2022) combined unimodal contrastive learning, masked volume in-painting, and image rotation prediction to learn 3D structural information. Chen et al. (2023) explored 3D masked image modeling, demonstrating faster convergence compared to simple contrastive methods (SimCLR and MAE). Valanarasu et al. (2023) used a reconstruction loss to restore 3D CT volumes from tokens corrupted by noise, downsampling, and local masking. Finally, Blankemeier et al. (2024) introduced Merlin, a dual-objective framework that leverages EHR data by first using contrastive learning to align radiology reports with CT volumes, followed by training on disease phenotype classification using labels derived from diagnostic codes recorded during contemporaneous hospital visits.

Current SSL approaches excel at capturing structural features in medical images, such as organ morphology for image segmentation. However, these learned features are static and largely fail to identify dynamic patterns and biomarkers that are predictive of future health risks. Our contributions help bridge this gap in several key ways by leveraging future medical events to guide pretraining, resulting in a more powerful image encoder for outcome prediction tasks. First, we employ a TTE objective to capture future patient health dynamics while accounting for right censoring, leverag-

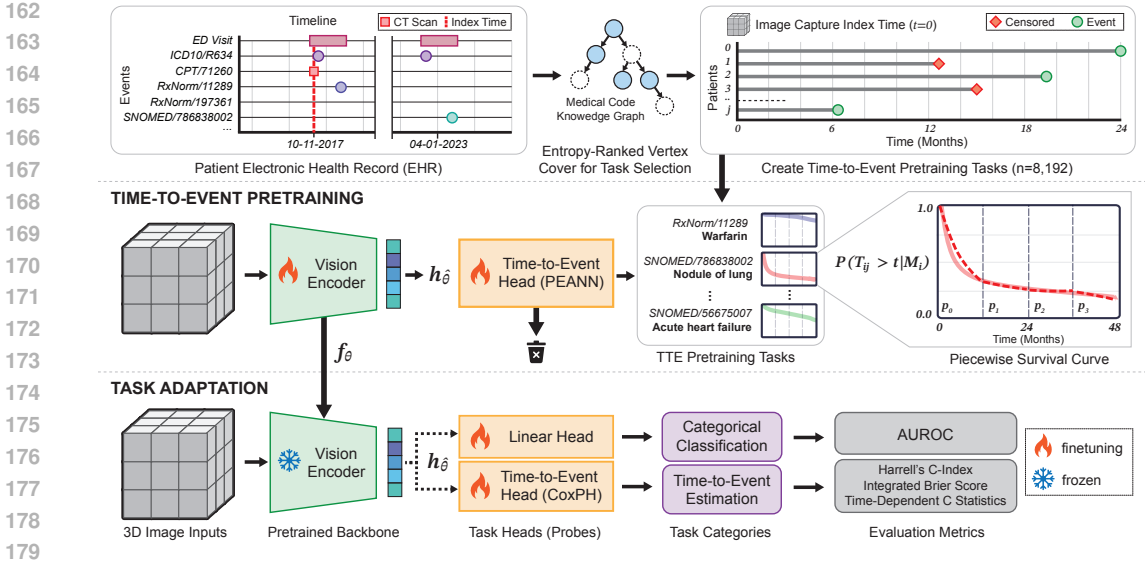


Figure 2: Overview of the proposed *time-to-event pretraining* pipeline. Patients’ longitudinal EHR timelines are transformed into large-scale, time-to-event (TTE) pretraining tasks. These tasks, which reflect informative temporal patterns for medical outcome prediction, are then used for continued pretraining (full fine-tuning) of a 3D vision encoder. The resulting encoder is then frozen and adapted to downstream tasks via different task heads for classification or TTE estimation.

ing longitudinal EHR data to greatly expand the scale, diversity, and temporal scope of pretraining supervision. Second, we evaluate pretraining using native 3D imaging architectures which are underexplored in prior TTE work. Finally, we evaluate the impact of TTE modeling in imaging at a larger scale than prior work, utilizing 18,945 CT scans—equivalent to 4.2 million 2D images—and 8,192 unique TTE tasks derived from longitudinal EHRs.

### 3 PRELIMINARIES

**Time-to-Event Modeling.** The objective of time-to-event modeling is to estimate the distribution of times  $T$  until an event of interest occurs. Observable data is denoted as  $\mathcal{O} = \{(\tilde{T}_i, \Delta_i, X_i^T) : i = 1, \dots, n\}$ , where  $X_i$  are the features of observation  $i$  and  $\Delta_i = \mathbb{I}(T_i \leq C_i)$  is an event indicator function whose value is 1 when the actual event time is observed. One complexity is that medical data is often right-censored, where the survival time  $T$  is not observed due to a loss of followup. With right-censoring, we do not observe  $T$  and instead observe time  $\tilde{T} = \min(T, C)$ , where the  $C$  is the censoring time. When  $\Delta_i = 0$ , we do not know the true survival time, but we know that it is greater than  $\tilde{T}_i$ . The mixture of known survival times and censored survival times necessitates methods that can estimate  $T$  in an unbiased manner despite the censorship.

Different time-to-event models use different definitions to model the instantaneous hazard  $\lambda(t)$ , where  $t$  represents the continuous time at which the hazard rate is evaluated. We use a piecewise exponential function (Kitchin et al., 1983) that splits the timeline into  $P$  distinct intervals, or *pieces*, each with a constant hazard rate  $\lambda(t)$ . Each piece  $p(t)$  has  $\lambda_{ip}(t) = C_{ip}$ , where  $C$  depends on the patient’s CT image  $i$  and the specific interval  $p$ .

**Piecewise Exponential Neural Network.** To make use of deep learning models, we use a neural network to define  $C_{ip}$ . Using a piecewise exponential neural network (PEANN) (Fornili et al., 2014) greatly simplifies large-scale pretraining compared to Cox PH-based methods, which require creating batches where samples are paired with at least one uncensored patient (Harrell Jr et al., 1996). We first use a neural network to derive a CT image representation  $M_{ip}$  for image  $i$ ’s information at time piece  $p$ . We then apply a linear layer followed by an exponential function to define hazard  $\lambda_{ip} = C_{ip} = \exp(A * M_{ip} + b)$  where  $A$  and  $b$  are learned parameters. With the hazard calculated,

we can then input it into the survival function:

$$S_i(t) = \prod_{p=1}^P \exp[-\lambda_{ip}(\min(t, E_p) - S_p)\mathbb{I}(t \geq S_p)] \quad (1)$$

where each time piece  $p$  has a starting point  $S_p$  and end point  $E_p$  and  $\mathbb{I}(\cdot)$  is an indicator function. The standard survival likelihood loss function for image  $i$  is then:

$$\mathcal{L}_i = [S_i(t)]^{1-\Delta_i} [f_i(t)]^{\Delta_i} \quad (2)$$

where  $\Delta_i$  is event indicator whose value is 1 when the actual future event is observed for image  $i$  and  $f_i$  is the probability density function  $f_i = -\frac{\partial S_i}{\partial t}$ .

A full derivation of the loss function, including calculating the derivative and plugging in the survival function, can be found in Appendix L.

## 4 TIME-TO-EVENT PRETRAINING

In this work, we are interested in training a 3D image encoder to learn representations optimized for estimating the distribution of event times for future clinical outcomes. Using years of follow-up EHR data after image capture, we generate time-to-event pretraining tasks that provide large-scale training signal for estimating this distribution. We outline our approach for time-to-event pretraining in Figure 2. Given the computational costs of pretraining 3D image architectures from scratch, we evaluate the benefits of TTE supervision via continued pretraining of an existing neural network  $f_\theta$ , where  $\theta$  represents the parameters of the pretrained backbone.

**Creating TTE Pretraining Tasks** EHR data captures a vast amount of structured information on patient demographics, diagnoses, procedures, medications, medical devices, social determinants of health, and other aspects of medical care. These data are encoded as timestamped, standardized identifiers called *medical codes* that map to ontologies (e.g., ICD10, RxNorm, CPT). These ontologies collectively represent a knowledge graph in the form of a directed acyclic graph (DAG). By treating each code as a separate task and organizing EHR data chronologically, we generate extensive training signals to model longitudinal health trajectories. This approach enables tasks such as predicting when a patient might develop lung cancer or be prescribed warfarin (a blood thinner used to prevent and treat blood clots), represented as time-to-event distributions conditioned on image features. However, naively treating each medical code as a task leads to millions of candidate pretraining tasks (4.3 million in our dataset). Many of these tasks are low-frequency or otherwise redundant given an ontology structure. To select pretraining tasks, we follow Steinberg et al. (2024)’s proposed conditional Shannon entropy selection measure, which treats ontology-aware task selection as a vertex cover problem (West et al., 2001). We select a set of medical codes that maximizes conditional entropy given a task budget, ontology DAG, and frequency distribution of observed medical codes.

We define our TTE pretraining procedure by predicting the time until the first occurrence (if there are multiple recurring ones) of a medical code as defined above. The survival function in each time piece is modeled from the time piece’s starting time, where the first time piece’s starting time is the index time (here the time of the CT scan exam). When there is no event in a time piece the loss will be 0. Each TTE task ( $n=8,192$ ) is modeled independently. We also apply censorship at patient death, which is the only competing risk. Given TTE labels for input image  $X_i$ , we use the loss described in Eq. 2 for full fine-tuning of  $f_\theta$ .

**Task Adaptation.** Once TTE pretraining is complete, we use the frozen encoder  $f_\theta$  to generate feature embeddings for each input image  $X_i$ , such that  $Z_i = f_\theta(X_i)$ . These embeddings are then passed to a task-specific head  $h_{\hat{\theta}}$ , which can either be a classification head or a time-to-event head, depending on the task. We found that a CoxPH task head (DeepSurv), which directly optimizes for Harrell’s C-index, performed better in practice than fine-tuning a PEANN task head, thus we used CoxPH for all TTE task adaptation. The classification head outputs prediction probabilities  $p_i(\hat{y}|Z_i)$  for discriminative tasks, while the survival head produces a time-dependent hazard score  $H_i(t|Z_i)$  for time-to-event (TTE) tasks. The model’s outputs are evaluated using task-appropriate metrics, e.g., AUROC for classification or Harrell’s C-index and time-dependent C-statistics for TTE tasks.

## 270 5 EXPERIMENTS

271  
272 **Hypotheses.** Our experiments measure the impact of TTE pretraining on an image encoder’s ability to generate representations useful for medical prognosis. We explore the following hypotheses:

- 273 1. TTE supervision improves data efficiency by utilizing temporal information from future EHRs and censored patients, increasing available task labels per-pretraining instance.
- 274 2. Existing supervised and SSL pretraining methods struggle to learn strong pixel biomarkers for disease risk due to missing temporal connections with pathologies that will be detected in the future. TTE supervision, in contrast, provides a simple approach to leveraging complex temporal information found in longitudinal EHRs for purposes of learning prognostic pixel biomarkers.
- 275 3. TTE supervision, when conducted as post hoc, continued pretraining, does not negatively impact performance on standard categorical classification tasks as used for diagnostic image labeling.

276 **Setup.** All image and EHR preprocessing details are outlined in Appendix A. For our PEANN, we use 8 piecewise time bins and 8,192 pretraining tasks (see Appendices N,O) per the best-performing hyperparameters from Steinberg et al. (2024). Each time bin is created by uniformly dividing the range between the cohort’s earliest and latest timestamps. We utilized two compute nodes from an on-premise cluster, each with 24 Intel Xeon 2.7GHz CPU cores, 200 GB RAM and 4 Nvidia A100 GPUs or 4 H100 GPUs.

277 **Datasets & Evaluation Tasks.** We use two 3D medical imaging datasets (see Table 1). INSPECT (Huang et al., 2024) is a multimodal dataset of paired CT scans and radiology notes where each patient is linked to their longitudinal EHR. This provides an average of 5 years follow-up data post-CT scan and, in aggregate, contains 225 million medical events. RSPECT (Colak et al., 2021) is an image-only dataset of CT scans annotated by radiologists for imaging biomarkers related to pulmonary embolism and cardiac function. We evaluate 3 task categories in this work:

- 304 • *Prognostic TTE:* Estimate the distribution of event times for a specific outcome (e.g, mortality). We predict the first occurrence of an event, as the first occurrence (or only occurrence in cases such as mortality) is of higher clinical utility than subsequent events.
- 305 • *Prognostic Classification:* Binarized classification formulation of TTE tasks using bucketed time bins. We use 1, 6, and 12 month bins. Unlike in TTE, censored patients are excluded in this category.
- 306 • *Diagnostic Classification:* Standard classification using diagnostic image label categories.

307  
308 INSPECT defines 3 prognostic binary tasks: hospital mortality, hospital readmission, and pulmonary hypertension (PH); and 1 binary diagnostic classification task (pulmonary embolism). Using the provided EHR data we define 5 additional prognostic TTE tasks for lung pathologies: ATX (Atelectasis), CMG (Cardiomegaly), CONS (Consolidation) EDM (Edema), and PEFF (Pleural Effusion). These tasks were selected based on their use in common chest medical imaging datasets (Irvin et al., 2019). Appendix B details how labels were assigned. RSPECT provides whole-volume labels for 9 diagnostic tasks, but no longitudinal outcome data. RSPECT does not provide a public test set, so we impose a 80/5/15 split on train (n=7,279) for our experiments.

309  
310  
311 **Architectures.** We evaluate three model architectures: SwinUNETR, DenseNet, and ResNet. SwinUNETR (Tang et al., 2022) was originally designed for medical image segmentation tasks, combining elements from the Swin Transformer and the UNETR (U-Net Transformer) architectures. Weights are learned using a reconstruction loss on 10,050 3D brain/chest CT and MRI data

	INSPECT	RSPECT
# Patients	19,402	7,279
# Train	18,945	5,823
# Valid	1,089	364
# Test	3,214	1,092
Imaging	✓	✓
EHR	✓	✗
TTE Tasks	✓	✗
Diag. Tasks	✓	✓
Scan Type	Chest CT	Chest CT

Table 1: Dataset summary statistics.

(Valanarasu et al., 2023), a much larger pretraining dataset than used by other public 3D models, e.g., Wasserthal et al. (2023). Following Merlin, we adapt DenseNet-121 (Huang et al., 2016) and ResNet-152 (He et al., 2015) by inflating their 2D pretrained ImageNet weights (specifically the filters and pooling kernels) as described in Carreira & Zisserman (2017). Note that ResNet and DenseNet parameters were initialized using 2D weight inflation, but the architecture is fully 3D, using 3D convolutions in each Res-block or Dense-block, different from 2.5D methods (Hung et al., 2024). This process enables us to input 3D CT images into the models for training. We also evaluated Merlin’s pretrained ResNet-152 backbone, but found it performed similar to our base ResNet-152 (see Appendix D), thus we use the base weight-inflated model for consistency across experiments.

**Model Baselines.** We evaluate the following continued pretraining approaches on all architectures:

- **base:** Baseline performance of the 3D pretrained SwinUNETR, DenseNet-121, and ResNet-152 models without continued pretraining on INSPECT. See Appendix G for a summary of the source pretraining datasets.
- **base/MTL:** Continued pretraining of the base model via multitask, supervised learning using the 8 INSPECT evaluation task labels. This controls for exposure to our training dataset in a consistent manner across architectures.
- **base/visit:** Continued pretraining using the same 8,192 tasks used for TTE supervision, but restricting label assignment to the same visit as the CT scan. This ablates the TTE component to learn temporal information and aligns more closely with the EHR-based supervision used by Merlin.
- **base/TTE:** Continued pretraining using TTE labels for 8,192 tasks occurring after the CT scan index timestamp.

All models except base use the same INSPECT training set examples for continued pretraining. Pretraining task labels are assigned per-CT scan and vary in density based on pretraining approach (Figure 3). Note that TTE supervision enables leveraging a patient’s entire future EHR, providing 3 times more training labels on average per CT-scan over per-visit labels. TTE also captures temporal structure and time-varying disease risk, providing supervision signal that is absent in predominant SSL methods for imaging.

For evaluating our pretrained encoder, we use the frozen encoder and lightweight task head strategy outlined in Section 4. For classification tasks, we use logistic regression as the classification head (linear probe) and for TTE tasks, we employ DeepSurv (Katzman et al., 2018) as the survival head. All model search hyperparameters for pretraining and adaptation are in Appendix E.

**Metrics.** We evaluate discrimination performance using AUROC for time-thresholded binary classification tasks and Harrell’s C-index (Harrell et al., 1982) for TTE tasks. Harrell’s C-index is a type of C statistic that summarizes the ability of a predictive model to rank patients. Harrell’s C-index requires the predictive model to output a single risk per patient, as opposed to risk over time, thus suitable for our DeepSurv head evaluation. We use the standard formulation for Harrell’s C-index, where it first finds all possible pairs  $P$  and one patient is known to have the event before another, then splits that set into  $P_{\text{correct}}$  (the higher predicted risk patient has the event first),  $P_{\text{tied}}$  (both patients are tied for the predicted risk), and  $P_{\text{incorrect}}$  (lower predicted risk patient actually has the event first):  $C_H = \frac{P_{\text{correct}} + 0.5 \times P_{\text{tied}}}{P_{\text{correct}} + P_{\text{tied}} + P_{\text{incorrect}}}$ . Appendix C contains additional TTE evaluation metrics for time-dependent C-statistics and the integrated Brier score.

All performance results and 95% confidence intervals are reported using a test set bootstrap of  $n = 1000$  replicates. Statistical significance was computed using a two-tailed Z-test (p-value at 0.05 for rejecting the null hypothesis that the difference between two sample means is zero). A complete set of statistical tests are in Appendix J.




Figure 3: Label density CDF by pretraining approach.

**Additional Experiments.** See the appendix for further experimental ablations including: subgroup performance (Appendix Q), task head capacity (Appendix K), and full fine-tuning vs. frozen backbone adaptation (Appendix I).

## 6 RESULTS

**Evaluating Prognostic Performance.** Table 2 reports performance for the prognostic binary classification formulations of outcome prediction for the original INSPECT tasks. This provides a simplified view of high and low-risk patients across time. We find that TTE pretrained models outperform all baselines across all architectures in our experiments. Here TTE pretraining provides an average of 22.6% performance over the base pretrained model and 15.4% average increase over base/visit. The base/MTL baseline also outperforms the base model, but underperforms TTE, highlighting the benefits of increasing pretraining tasks. Comparing base/TTE versus base/visit is more informative, as both approaches use the same number of pretraining tasks (8,192) but base/TTE substantially improves the density of the label per image during pretraining (Figure 3). Table 3 reports performance of all original INSPECT tasks and our five new outcomes using Harrell’s C-index. Here TTE pretraining largely outperforms all of our baselines across all architectures.

Model	Mortality			Readmission			PH
	1M	6M	12M	1M	6M	12M	12M
SwinUNETR <sub>base</sub>	0.693	0.684	0.685	0.507	0.538	0.569	0.597
SwinUNETR <sub>base/MTL</sub>	0.676	0.700	0.697	0.502	0.543	0.551	0.606
SwinUNETR <sub>base/visit</sub>	0.693	0.716	0.670	<u>0.560</u>	0.554	0.528	0.560
SwinUNETR <sub>base/TTE</sub>	<b>0.827</b>	<b>0.808</b>	<b>0.788</b>	<b>0.582</b>	<b>0.612</b>	<b>0.607</b>	<b>0.672</b>
DenseNet-121 <sub>base</sub>	0.616	0.568	0.575	0.512	0.532	0.524	0.506
DenseNet-121 <sub>base/MTL</sub>	0.665	0.596	0.602	0.533	0.557	0.547	0.577
DenseNet-121 <sub>base/visit</sub>	0.649	0.698	0.692	0.504	0.538	0.567	0.599
DenseNet-121 <sub>base/TTE</sub>	<b>0.770</b>	<b>0.730</b>	<b>0.725</b>	<b>0.629</b>	<b>0.643</b>	<b>0.637</b>	<b>0.689</b>
ResNet-152 <sub>base</sub>	0.583	0.557	0.554	0.536	0.537	0.509	0.537
ResNet-152 <sub>base/MTL</sub>	0.726	0.715	0.647	<u>0.563</u>	0.564	0.566	0.570
ResNet-152 <sub>base/visit</sub>	0.691	0.636	0.643	0.562	0.564	0.566	0.567
ResNet-152 <sub>base/TTE</sub>	<b>0.804</b>	<b>0.798</b>	<b>0.792</b>	<b>0.602</b>	<b>0.649</b>	<b>0.657</b>	<b>0.718</b>

Table 2: Prognostic binary classification performance for INSPECT tasks on Logistic Regression (1, 6, 12 month time horizon bins) reported as the mean AUROC of a test set bootstrap (n=1000). **Bold** indicates the best performer. Underlined indicates no statistically significant difference versus the \*<sub>base/TTE</sub> models.

**Evaluating Diagnostic Performance.** We are also interested in assessing how TTE continued pretraining may impact standard image classification tasks, corresponding to diagnostic labeling of medical imaging for current disease biomarkers. Table 4 outlines 8 image biomarker tasks. Here, performance of all 3D model architectures is poor, especially base models. For almost all diagnostic tasks, TTE pretraining performs the same (i.e., statistically indistinguishable, detailed numbers shown in Table 17) as all other tested pretraining approaches. This aligns with the intuition that visit-level labels reflect current clinical events, should encode the same level of diagnostic information as TTE pretraining. This also aligns with the small performance gains reported in prior work when tasks derived from EHR codes to supervise image models (Blankemeier et al., 2024). Note that while these tasks reflect observable pixel biomarkers present in images, there is also overlap with pixel biomarkers indicative of future risk. For example, a RV/LV ratio  $\geq 1$ , defined as the ratio of the right ventricular (RV) diameter to the left ventricular (LV) diameter, is indicative of increased risk of mortality (Lu et al., 2012). Here TTE pretraining yields statistically significant performance improvements across all architectures and pretraining methods.



Model	Harrell's C-Index $\uparrow$							
	Mort.	Readm.	PH	ATX	CMG	CONS	EDM	PEFF
SwinUNETR <sub>base</sub>	0.717	0.653	0.696	0.558	0.662	0.549	0.697	0.641
SwinUNETR <sub>base/MTL</sub>	0.672	0.671	0.665	0.681	0.677	0.668	0.715	0.668
SwinUNETR <sub>base/visit</sub>	0.671	<u>0.716</u>	0.697	0.719	0.717	0.718	<u>0.717</u>	0.714
SwinUNETR <sub>base/TTE</sub>	<b>0.738</b>	<b>0.723</b>	<b>0.724</b>	<b>0.739</b>	<b>0.739</b>	<b>0.738</b>	<b>0.738</b>	<b>0.738</b>
DenseNet-121 <sub>base</sub>	0.505	0.505	0.541	0.505	0.505	0.506	0.505	0.536
DenseNet-121 <sub>base/MTL</sub>	0.589	0.590	0.591	0.593	0.589	0.586	0.587	0.590
DenseNet-121 <sub>base/visit</sub>	0.675	0.675	<u>0.699</u>	0.675	0.662	0.676	0.676	0.675
DenseNet-121 <sub>base/TTE</sub>	<b>0.732</b>	<b>0.723</b>	<b>0.726</b>	<b>0.720</b>	<b>0.711</b>	<b>0.712</b>	<b>0.725</b>	<b>0.723</b>
ResNet-152 <sub>base</sub>	0.505	0.560	0.505	0.577	0.505	0.559	0.505	0.536
ResNet-152 <sub>base/MTL</sub>	<u>0.701</u>	0.686	0.656	0.663	0.562	0.656	0.660	0.572
ResNet-152 <sub>base/visit</sub>	0.656	0.702	0.643	0.703	0.705	0.703	0.700	<u>0.716</u>
ResNet-152 <sub>base/TTE</sub>	<b>0.732</b>	<b>0.739</b>	<b>0.735</b>	<b>0.728</b>	<b>0.727</b>	<b>0.727</b>	<b>0.737</b>	<b>0.737</b>

Table 3: Prognostic TTE performance for INSPECT, measured by Harrell’s C-Index. **Bold** indicates the best performance. Underlined indicates no statistically significant difference versus the \*<sub>base/TTE</sub> models.

Model	PE						RV/LV Ratio	
	Left	Cent.	Right	Chronic	Acute	Indet.	<1	$\geq 1$
SwinUNETR <sub>base</sub>	0.573	0.571	0.573	<u>0.507</u>	0.581	0.721	0.525	0.578
SwinUNETR <sub>base/MTL</sub>	0.583	<u>0.633</u>	0.591	<u>0.504</u>	0.594	<u>0.732</u>	0.517	0.598
SwinUNETR <sub>base/visit</sub>	0.545	<b>0.664</b>	<u>0.562</u>	<b>0.556</b>	<u>0.595</u>	<u>0.741</u>	0.521	0.546
SwinUNETR <sub>base/TTE</sub>	<b>0.634</b>	0.651	<b>0.633</b>	0.525	<b>0.716</b>	<b>0.781</b>	<b>0.643</b>	<b>0.636</b>
DenseNet-121 <sub>base</sub>	0.596	0.615	0.596	<u>0.565</u>	0.676	0.724	<u>0.581</u>	0.607
DenseNet-121 <sub>base/MTL</sub>	0.612	0.638	0.615	0.504	<u>0.664</u>	0.713	0.594	0.636
DenseNet-121 <sub>base/visit</sub>	0.595	0.633	0.605	<u>0.571</u>	<b>0.677</b>	<u>0.748</u>	0.573	0.615
DenseNet-121 <sub>base/TTE</sub>	<b>0.647</b>	<b>0.716</b>	<b>0.644</b>	<b>0.586</b>	0.665	<b>0.762</b>	<b>0.623</b>	<b>0.686</b>
ResNet-152 <sub>base</sub>	<u>0.654</u>	<u>0.687</u>	<u>0.618</u>	0.507	0.695	0.704	<u>0.593</u>	0.626
ResNet-152 <sub>base/MTL</sub>	<u>0.658</u>	<b>0.693</b>	<u>0.621</u>	<u>0.506</u>	0.685	0.703	<u>0.587</u>	0.642
ResNet-152 <sub>base/visit</sub>	0.646	0.665	0.625	<b>0.548</b>	<b>0.743</b>	<u>0.718</u>	0.566	0.665
ResNet-152 <sub>base/TTE</sub>	<b>0.657</b>	0.687	<b>0.642</b>	0.504	0.588	<b>0.722</b>	<b>0.597</b>	<b>0.708</b>

Table 4: Diagnostic binary classification performance for RSPECT, reported as the mean AUROC of a test set bootstrap (n=1000). **Bold** indicates the best performer. Underlined indicates no statistically significant difference versus the \*<sub>base/TTE</sub> models.

## 7 DISCUSSION AND CONCLUSION

Medical images hold significant, untapped potential as sources of imaging biomarkers to predict future disease risk. However, current SSL approaches for imaging largely fail to capture the temporal dynamics of long-term disease progression, inherently capturing only static, structural information. Building on this observation, our work explores a TTE pretraining technique that directly incorporates future temporal information at scale, yielding several insights.

**Current Pretraining Struggles to Learn Prognostic Pixel Biomarkers.** Existing supervised and SSL pretraining methods consistently showed lower AUROC across prognostic tasks in our experiments. When compared against base models, TTE pretraining showed an average increase in AUROC of 0.128 (95% CI: [0.075, 0.158]), Harrell’s C-index 0.166 (95% CI [0.138, 0.490]) over three architectures on prognostic binary classification and prognostic TTE tasks respectively. This suggests that off-the-shelf pretraining methods struggle to learn pixel biomarkers associated with future disease risk, likely due to missing temporal links to future pathologies. This underscores Huang

et al. (2020)’s findings that predicting disease prognosis with high accuracy and certainty remains a challenging task. Therefore, incorporating explicit temporal supervision into the pretraining process may be essential for improving prognostic tasks’ performance.

**TTE Pretraining Improves Prognostic Performance.** When compared against base/MTL and base/visit, TTE supervision improves performance on Harrell’s C-index on prognostic TTE task tasks by 0.093 (95% CI [0.086, 0.554]) and, 0.038 (95% CI [0.017, 0.370]) respectively. Additionally TTE supervision does not negatively impact diagnostic binary classification for significant difference between base and TTE across RSPECT tasks, shown in Table 17.

**TTE Supervision Improves Training Data Efficiency.** By leveraging temporal information from the future in EHRs and censored patients, TTE pretraining increases label density to boost AUROC for prognostic binary classification tasks by 0.093 (95% CI [0.053, 0.134]) and 0.092 (95% CI [0.051, 0.135]) when compared against base/MTL and base/visit models respectively. TTE supervision increases available task labels per-training instance by 3x on average (Figure 3). This underscores the potential of TTE pretraining as a viable objective for scaling medical imaging AI, given that expert-level annotation is time-consuming, expensive, and difficult to collect (Dgani et al., 2018; Tajbakhsh et al., 2021; Aljabri et al., 2022). Increasing the data efficiency of a given training example also contributes to reducing compute costs.

**Limitations.** First, our study focuses exclusively on evaluating 3D vision encoders, which demand significant memory and computational resources—80GB memory GPUs for SwinUNETR and 40GB memory GPUs for DenseNet/ResNet architectures. Second, while our study surpasses previous work in imaging-based TTE in terms of scale, our pretraining dataset remains relatively small compared to modern, general-purpose datasets (Schuhmann et al., 2022) and recently released medical datasets (Xie et al., 2024). Furthermore, we focus exclusively on a single modality, CT scans. Expanding both the scale and diversity of the pretraining data mixtures (e.g., other imaging modalities, historical EHR and clinical text) could enhance performance or lead to a deeper understanding of the trade-offs between different architectures. Finally, since this work focuses on evaluating encoder quality, we only evaluated frozen encoders with smaller, lightweight, supervised task heads for adaptation. Alternative adaption methods under different sample assumptions, e.g., zero/few-shot learning, may reveal different performance trade-offs.

**Conclusion.** This work presents the first empirical study of using time-to-event pretraining for 3D medical vision models. By leveraging longitudinal EHRs and defining a time-to-event pretraining objective (comprising 8,192 tasks), we embed long term outcome information into the image model during pretraining. Doing so results in an average threefold increase in training labels compared to limiting just to a patient’s current EHR visit. We observe substantial improvements in prediction performance for future events, with increases of up to a 31.6% in AUROC and a 40.5% improvement in Harrell’s C-index, without negatively impacting standard binary classification tasks (e.g., image labeling for diagnostic tasks). Our results reveal a clear need for pretraining datasets for 3D medical foundation models to include tasks that capture long-term temporal structure, demonstrated here through our use of time-to-event supervision. This study on the utility of using longitudinal EHR records as a supervision source for future-guided pretraining of 3D medical imaging models lays the groundwork for innovative ways to combine EHR and imaging modalities for clinical risk prediction.

## REFERENCES

- Asad Aali, Andrew Johnston, Louis Blankemeier, Dave Van Veen, Laura T Derry, David Svec, Jason Hom, Robert D Boutin, and Akshay S Chaudhari. Detecting underdiagnosed medical conditions with deep learning-based opportunistic ct imaging. [arXiv preprint arXiv:2409.11686](#), 2024.
- Manar Aljabri, Manal AlAmir, Manal AlGhamdi, Mohamed Abdel-Mottaleb, and Fernando Collado-Mesa. Towards a better understanding of annotation tools for medical imaging: a survey. *Multimedia tools and applications*, 81(18):25877–25911, 2022.
- Shekoofeh Azizi, Laura Culp, Jan Freyberg, Basil Mustafa, Sebastien Baur, Simon Kornblith, Ting Chen, Patricia MacWilliams, S Sara Mahdavi, Ellery Wulczyn, et al. Robust and efficient medical imaging with self-supervision. [arXiv preprint arXiv:2205.09723](#), 2022.

- 540 Kaustav Bera, Nathaniel Braman, Amit Gupta, Vamsidhar Velcheti, and Anant Madabhushi. Pre-  
541 dicting cancer outcomes with radiomics and artificial intelligence in radiology. Nature reviews  
542 Clinical oncology, 19(2):132–146, 2022.
- 543
- 544 Louis Blankemeier, Joseph Paul Cohen, Ashwin Kumar, Dave Van Veen, Syed Jamal Safdar  
545 Gardezi, Magdalini Paschali, Zhihong Chen, Jean-Benoit Delbrouck, Eduardo Reis, Cesar Truys,  
546 et al. Merlin: A vision language foundation model for 3d computed tomography. arXiv preprint  
547 arXiv:2406.06512, 2024.
- 548 Patrick Breheny. Accelerated failure time models. University Lecture, 2015.
- 549
- 550 Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics  
551 dataset. In proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp.  
552 6299–6308, 2017.
- 553
- 554 Trenton Chang, Michael W Sjoding, and Jenna Wiens. Disparate censorship & undertesting: A  
555 source of label bias in clinical machine learning. In Machine Learning for Healthcare Conference,  
556 pp. 343–390. PMLR, 2022.
- 557
- 558 Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for  
559 contrastive learning of visual representations. In International conference on machine learning,  
560 pp. 1597–1607. PMLR, 2020.
- 561
- 562 Zekai Chen, Devansh Agarwal, Kshitij Aggarwal, Wiem Safta, Mariann Micsinai Balan, and Kevin  
563 Brown. Masked image modeling advances 3d medical image analysis. In Proceedings of the  
564 IEEE/CVF Winter Conference on Applications of Computer Vision, pp. 1970–1980, 2023.
- 565
- 566 Errol Colak, Felipe C Kitamura, Stephen B Hobbs, Carol C Wu, Matthew P Lungren, Luciano M  
567 Prevedello, Jayashree Kalpathy-Cramer, Robyn L Ball, George Shih, Anouk Stein, et al. The rsna  
568 pulmonary embolism ct dataset. Radiology: Artificial Intelligence, 3(2):e200254, 2021.
- 569
- 570 David Collett. Modelling survival data in medical research. Chapman and Hall/CRC, 2023.
- 571
- 572 David R Cox. Regression models and life-tables. Journal of the Royal Statistical Society: Series B  
573 (Methodological), 34(2):187–202, 1972.
- 574
- 575 Yair Dgani, Hayit Greenspan, and Jacob Goldberger. Training a neural network based on unreliable  
576 human annotation of medical images. In 2018 IEEE 15th International symposium on biomedical  
577 imaging (ISBI 2018), pp. 39–42. IEEE, 2018.
- 578
- 579 Marco Fornili, Federico Ambrogi, Patrizia Boracchi, and Elia Biganzoli. Piecewise exponential  
580 artificial neural networks (peann) for modeling hazard function with right censored data. In  
581 Computation Intelligence Methods for Bioinformatics and Biostatistics, pp. 125–136, 07 2014.  
582 ISBN 978-3-319-09041-2. doi: 10.1007/978-3-319-09042-9\_9.
- 583
- 584 Erika Graf, Claudia Schmoor, Willi Sauerbrei, and Martin Schumacher. Assessment and comparison  
585 of prognostic classification schemes for survival data. Statistics in medicine, 18(17-18):2529–  
586 2545, 1999.
- 587
- 588 Frank E Harrell, Robert M Califf, David B Pryor, Kerry L Lee, and Robert A Rosati. Evaluating the  
589 yield of medical tests. Jama, 247(18):2543–2546, 1982.
- 590
- 591 Frank E Harrell Jr, Kerry L Lee, and Daniel B Mark. Multivariable prognostic models: issues in  
592 developing models, evaluating assumptions and adequacy, and measuring and reducing errors.  
593 Statistics in medicine, 15(4):361–387, 1996.
- 594
- 595 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recog-  
596 nition. arxiv e-prints. arXiv preprint arXiv:1512.03385, 10, 2015.
- 597
- 598 Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked au-  
599 toencoders are scalable vision learners. In Proceedings of the IEEE/CVF conference on computer  
600 vision and pattern recognition, pp. 16000–16009, 2022.

- 594 Patrick J Heagerty and Yingye Zheng. Survival model predictive accuracy and roc curves.  
595 Biometrics, 61(1):92–105, 2005.
- 596
- 597 Gregory Holste, Mingquan Lin, Ruiwen Zhou, Fei Wang, Lei Liu, Qi Yan, Sarah H Van Tassel, Kyle  
598 Kovacs, Emily Y Chew, Zhiyong Lu, et al. Harnessing the power of longitudinal medical imaging  
599 for eye disease prognosis using transformer-based sequence modeling. NPJ Digital Medicine, 7  
600 (1):216, 2024.
- 601 Gao Huang, Zhuang Liu, and Kilian Q Weinberger. Densely connected convolutional networks.  
602 corr. arXiv preprint arXiv:1608.06993, 2016.
- 603
- 604 Shigao Huang, Jie Yang, Simon Fong, and Qi Zhao. Artificial intelligence in cancer diagnosis and  
605 prognosis: Opportunities and challenges. Cancer letters, 471:61–71, 2020.
- 606
- 607 Shih-Cheng Huang, Zepeng Huo, Ethan Steinberg, Chia-Chun Chiang, Curtis Langlotz, Matthew  
608 Lungren, Serena Yeung, Nigam Shah, and Jason Fries. Inspect: A multimodal dataset for pa-  
609 tient outcome prediction of pulmonary embolisms. Advances in Neural Information Processing  
610 Systems, 36, 2024.
- 611 Alex Ling Yu Hung, Haoxin Zheng, Kai Zhao, Xiaoxi Du, Kaifeng Pang, Qi Miao, Steven S Raman,  
612 Demetri Terzopoulos, and Kyunghyun Sung. Csam: A 2.5 d cross-slice attention module for  
613 anisotropic volumetric medical image segmentation. In Proceedings of the IEEE/CVF Winter  
614 Conference on Applications of Computer Vision, pp. 5923–5932, 2024.
- 615
- 616 Jeremy Irvin, Pranav Rajpurkar, Michael Ko, Yifan Yu, Silviana Ciurea-Ilcus, Chris Chute, Hen-  
617 rik Marklund, Behzad Haghgoo, Robyn L. Ball, Katie S. Shpanskaya, Jayne Seekins, David A.  
618 Mong, Safwan S. Halabi, Jesse K. Sandberg, Ricky Jones, David B. Larson, Curtis P. Langlotz,  
619 Bhavik N. Patel, Matthew P. Lungren, and Andrew Y. Ng. Chexpert: A large chest radiograph  
620 dataset with uncertainty labels and expert comparison. CoRR, abs/1901.07031, 2019. URL  
621 <http://arxiv.org/abs/1901.07031>.
- 622 Hemant Ishwaran, Udaya B Kogalur, Eugene H Blackstone, and Michael S Lauer. Random survival  
623 forests. 2008.
- 624
- 625 Edward L Kaplan and Paul Meier. Nonparametric estimation from incomplete observations. Journal  
626 of the American statistical association, 53(282):457–481, 1958.
- 627
- 628 Jared L Katzman, Uri Shaham, Alexander Cloninger, Jonathan Bates, Tingting Jiang, and Yuval  
629 Kluger. Deepsurv: personalized treatment recommender system using a cox proportional hazards  
630 deep neural network. BMC medical research methodology, 18:1–12, 2018.
- 631 Alexander Ke, William Ellsworth, Oishi Banerjee, Andrew Y Ng, and Pranav Rajpurkar. Chextran-  
632 ser: performance and parameter efficiency of imagenet models for chest x-ray interpretation. In  
633 Proceedings of the conference on health, inference, and learning, pp. 116–124, 2021.
- 634
- 635 John Kitchin, Naftali A Langberg, and Frank Proschan. A new-method for estimating life distribu-  
636 tions from incomplete data. Statistics & Risk Modeling, 1(3):241–256, 1983.
- 637
- 638 David G Kleinbaum and Mitchel Klein. Survival analysis a self-learning text. Springer, 1996.
- 639
- 640 Adones Leslie, AJ Jones, and PR Goddard. The influence of clinical information on the reporting of  
641 ct by radiologists. The British journal of radiology, 73(874):1052–1055, 2000.
- 642
- 643 Lin Lu, Laurent Derclé, Binsheng Zhao, and Lawrence H Schwartz. Deep learning for the prediction  
644 of early on-treatment response in metastatic colorectal cancer from serial medical imaging. Nature  
communications, 12(1):6654, 2021.
- 645
- 646 Michael T Lu, Shadpour Demehri, Tianxi Cai, Layla Parast, Andetta R Hunsaker, Samuel Z Gold-  
647 haber, and Frank J Rybicki. Axial and reformatted four-chamber right ventricle-to-left ventricle  
diameter ratios on pulmonary ct angiography as predictors of death after acute pulmonary em-  
bolism. American Journal of Roentgenology, 198(6):1353–1360, 2012.

- 648 Elsa Negro-Calduch, Natasha Azzopardi-Muscat, Ramesh S Krishnamurthy, and David Novillo-  
649 Ortiz. Technological progress in electronic health record system optimization: Systematic review  
650 of systematic literature reviews. International journal of medical informatics, 152:104507, 2021.  
651
- 652 Ziad Obermeyer, Brian Powers, Christine Vogeli, and Sendhil Mullainathan. Dissecting racial bias  
653 in an algorithm used to manage the health of populations. Science, 366(6464):447–453, 2019.  
654
- 655 OHDSI. Omop common data model. <https://ohdsi.github.io/CommonDataModel/index.html>, 2023.  
656 Accessed: 2023-06-07.
- 657 Suraj Pai, Dennis Bontempi, Ibrahim Hadzic, Vasco Prudente, Mateo Sokač, Tafadzwa L Chaunzwa,  
658 Simon Bernatz, Ahmed Hosny, Raymond H Mak, Nicolai J Birkbak, et al. Foundation model for  
659 cancer imaging biomarkers. Nature machine intelligence, 6(3):354–367, 2024.  
660
- 661 Stephen R Pfohl, Heather Cole-Lewis, Rory Sayres, Darlene Neal, Mercy Asiedu, Awa Dieng, Ne-  
662 nad Tomasev, Qazi Mamunur Rashid, Shekoofeh Azizi, Negar Rostamzadeh, et al. A toolbox for  
663 surfacing health equity harms and biases in large language models. Nature Medicine, pp. 1–11,  
664 2024.
- 665 Luís Pinto-Coelho. How artificial intelligence is shaping medical imaging technology: A survey of  
666 innovations and applications. Bioengineering, 10(12):1435, 2023.  
667
- 668 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal,  
669 Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual  
670 models from natural language supervision. In International conference on machine learning, pp.  
671 8748–8763. PMLR, 2021.
- 672 Khaled Saab, Tao Tu, Wei-Hung Weng, Ryutarō Tanno, David Stutz, Ellery Wulczyn, Fan Zhang,  
673 Tim Strother, Chunjong Park, Elahe Vedadi, et al. Capabilities of gemini models in medicine.  
674 arXiv preprint arXiv:2404.18416, 2024.
- 675 Rinku Saikia and Manash Pratim Barman. A review on accelerated failure time models. Int J Stat  
676 Syst, 12(2):311–322, 2017.  
677
- 678 Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi  
679 Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An  
680 open large-scale dataset for training next generation image-text models. Advances in Neural  
681 Information Processing Systems, 35:25278–25294, 2022.
- 682 Bob Segert. Athena’s OMOP vocabulary list. [https://athena.ohdsi.org/vocabulary/  
683 list](https://athena.ohdsi.org/vocabulary/list), 2023. Accessed: 2023-08-16.  
684
- 685 Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh,  
686 and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based local-  
687 ization. In Proceedings of the IEEE international conference on computer vision, pp. 618–626,  
688 2017.
- 689 Michelle Shu, Richard Strong Bowen, Charles Herrmann, Gengmo Qi, Michele Santacatterina, and  
690 Ramin Zabih. Deep survival analysis with longitudinal x-rays for covid-19. In Proceedings of the  
691 IEEE/CVF International Conference on Computer Vision, pp. 4046–4055, 2021.  
692
- 693 Hari Sowrirajan, Jingbo Yang, Andrew Y Ng, and Pranav Rajpurkar. Moco pretraining improves  
694 representation and transferability of chest x-ray models. In Medical Imaging with Deep Learning,  
695 pp. 728–744. PMLR, 2021.
- 696 Harold C Sox, Michael C Higgins, Douglas K Owens, and Gillian Sanders Schmidler. Medical  
697 decision making. John Wiley & Sons, 2024.  
698
- 699 Anuroop Sriram, Matthew Muckley, Koustuv Sinha, Farah Shamout, Joelle Pineau, Krzysztof J  
700 Geras, Lea Azour, Yindalon Aphinyanaphongs, Nafissa Yakubova, and William Moore. Covid-19  
701 prognosis via self-supervised representation learning and multi-image prediction. arXiv preprint  
arXiv:2101.04909, 2021.

- 702 Ethan Steinberg, Jason Alan Fries, Yizhe Xu, and Nigam Shah. MOTOR: A time-to-event founda-  
703 tion model for structured medical records. In The Twelfth International Conference on Learning  
704 Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024. OpenReview.net, 2024. URL  
705 <https://openreview.net/forum?id=Nialiwi2V6>.  
706
- 707 Nima Tajbakhsh, Holger Roth, Demetri Terzopoulos, and Jianming Liang. Guest editorial  
708 annotation-efficient deep learning: the holy grail of medical imaging. IEEE transactions on  
709 medical imaging, 40(10):2526–2533, 2021.
- 710 Yucheng Tang, Dong Yang, Wenqi Li, Holger R Roth, Bennett Landman, Daguang Xu, Vishwesh  
711 Nath, and Ali Hatamizadeh. Self-supervised pre-training of swin transformers for 3d medical  
712 image analysis. In Proceedings of the IEEE/CVF conference on computer vision and pattern  
713 recognition, pp. 20730–20740, 2022.
- 714 Jeya Maria Jose Valanarasu, Yucheng Tang, Dong Yang, Ziyue Xu, Can Zhao, Wenqi Li, Vishal M  
715 Patel, Bennett Landman, Daguang Xu, Yufan He, et al. Disruptive autoencoders: Leveraging  
716 low-level features for 3d medical image pre-training. arXiv preprint arXiv:2307.16896, 2023.
- 717 Zifeng Wang, Zhenbang Wu, Dinesh Agarwal, and Jimeng Sun. Medclip: Contrastive learning from  
718 unpaired medical images and text. In Proceedings of the 2022 Conference on Empirical Methods  
719 in Natural Language Processing, pp. 3876–3887, 2022.
- 720 Jakob Wasserthal, Hanns-Christian Breit, Manfred T Meyer, Maurice Pradella, Daniel Hinck,  
721 Alexander W Sauter, Tobias Heye, Daniel T Boll, Joshy Cyriac, Shan Yang, et al. Totalseg-  
722 mentator: robust segmentation of 104 anatomic structures in ct images. Radiology: Artificial  
723 Intelligence, 5(5), 2023.
- 724 Douglas Brent West et al. Introduction to graph theory, volume 2. Prentice hall Upper Saddle River,  
725 2001.
- 726 Yiting Xie and David Richmond. Pre-training on grayscale imagenet improves medical image clas-  
727 sification. In Proceedings of the European conference on computer vision (ECCV) workshops,  
728 pp. 0–0, 2018.
- 729 Yunfei Xie, Ce Zhou, Lang Gao, Juncheng Wu, Xianhang Li, Hong-Yu Zhou, Sheng Liu, Lei Xing,  
730 James Zou, Cihang Xie, et al. Medtrinity-25m: A large-scale multimodal dataset with multigran-  
731 ular annotations for medicine. arXiv preprint arXiv:2408.02900, 2024.
- 732 Adam Yala and Kevin S Hughes. Rethinking risk modeling with machine learning. Annals of  
733 Surgical Oncology, 30(12):6950–6952, 2023.
- 734 Yuhao Zhang, Hang Jiang, Yasuhide Miura, Christopher D Manning, and Curtis P Langlotz. Con-  
735 trastive learning of medical visual representations from paired images and text. In Machine  
736 Learning for Healthcare Conference, pp. 2–25. PMLR, 2022.
- 737 Xinliang Zhu, Jiawen Yao, and Junzhou Huang. Deep convolutional neural network for survival  
738 analysis with pathological images. In 2016 IEEE international conference on bioinformatics and  
739 biomedicine (BIBM), pp. 544–547. IEEE, 2016.
- 740  
741  
742  
743  
744  
745  
746  
747  
748  
749  
750  
751  
752  
753  
754  
755

## ETHICS STATEMENT

Research involving de-identified publicly available data does not require Institutional Review Board (IRB) approval. However, to uphold standards and safeguard patient privacy, we follow the health-care machine learning reproducibility recommendations set by the Medical AI Research Foundations (Azizi et al., 2022). Below, we outline specific ethical considerations and include references to more detailed discussions in the appendix.

**Data Deidentification:** In accordance with applicable privacy laws and institutional guidelines, all data utilized for model training and evaluation has been de-identified by the original dataset authors. Both datasets used in this study, INSPECT (Huang et al., 2024) and RSPECT (Colak et al., 2021), underwent thorough removal of protected health information prior to their public release and were approved by an Institutional Review Board.

**Security, Data Storage, and Compliance:** All authors involved in data handling have completed institutional training on HIPAA and data privacy before engaging with the data. All training data was stored in a HIPAA-compliant compute environment.

**Algorithmic Bias** Healthcare machine learning models can be susceptible to algorithmic bias, leading to unfavorable outcomes for underrepresented subgroups (Obermeyer et al., 2019). Bias mitigation in medical foundation models remains an ongoing research challenge (Pfohl et al., 2024) and is not covered in this study. However, we take two steps to mitigate risk. First, all of our continued pretrained model releases includes a Data Use Agreement (DUA) that explicitly prohibits direct medical care. Second, in line with the recommendations from Chang et al. (2022), we conduct an analysis in Appendix Q to assess performance across sensitive subgroups, ensuring our pretraining technique does not unfairly disadvantage any group compared to existing methods. We evaluate performance using the AUROC (bootstrapped, n=1000) on 7 binary prognostic tasks, comparing TTE against base, showing that TTE pretraining does not reduce performance for sensitive groups and generally improves risk ranking across all groups.

## REPRODUCIBILITY STATEMENT

The code artifact necessary for reproducing the experiments in this paper can be found in the supplemental materials as a zip file. The anonymous Github link is [https://anonymous.4open.science/r/future\\_guided\\_pretraining-DA6C](https://anonymous.4open.science/r/future_guided_pretraining-DA6C). Our hyperparameter search grids can be found in Appendix E. All base pretraining weights are publicly available as detailed in Table 10. To ensure reproducibility, all experiments use researcher accessible, public medical datasets.

## APPENDIX

## A DATA PREPROCESSING

**CT Scans.** Each CT scan is preprocessed by extracting pixel data and applying a linear transformation to Hounsfield Units (HU) using the rescale slope and intercept values from the original DICOM records. To retain fine-grained details, we ensure axial slices have a thickness between 1 mm and 3 mm. Finally, we pad and center crop the images to 224 x 224 pixels.

**Electronic Health Records.** INSPECT’s EHR data is provided in the Observational Medical Outcomes Partnership (OMOP) Common Data Model (CDM) schema, a standardized framework that harmonizes healthcare data from various sources to support large-scale analysis (OHDSI, 2023). We use the Athena OHDSI Vocabularies Repository (Segert, 2023) (OMOP vocabulary version: v20240830) as our knowledge graph for generating tasks.

## B INSPECT NEW TTE TASK DEFINITIONS

We have selected a set of commonly used pulmonary disease tasks (Irvin et al., 2019), that are coded in INSPECT dataset’s EHR events. We use these common tasks to benchmark the time-to-event performance of our proposed TTE formulation. Each patient is assigned time-to-event=True after CT’s capture for pulmonary hypertension, atelectasis, cardiomegaly, consolidation, edema, and pleural effusion. These labels indicate either the time until the first occurrence of each condition or the time until the patient is censored. The description can be found in Figure 4 for diagnostic and prognostic labels.




Figure 4: Overview of Label Definitions: Diagnostic tasks use labels derived from the same hospital visit as the CT scan. Prognostic tasks involve future medical events from patients’ EHR timelines and are categorized into binary prognostic labels and time-to-event (TTE) prognostic labels. Note that for TTE tasks, only the time until the first occurrence is labeled.



## C ADDITIONAL TIME-TO-EVENT MODEL METRICS

We calculate additional metrics for time-to-event task modeling: the integrated Brier score (IBS) Graf et al. (1999), shown in Table 5, and time-dependent C-statistics Heagerty & Zheng (2005), shown in Table 6. IBS is intended to assess the discrepancy between predicted survival probabilities and observed outcomes. At each time point, the Brier score calculates the squared difference between the predicted survival probability and the actual event or censoring status. By integrating over the time range from the start to a specified maximum horizon, it summarizes the model’s predictive performance over the entire period of interest. For time-dependent C-statistics, we adopt the incident and dynamic definitions for time-dependent sensitivity and specificity Heagerty & Zheng (2005) to derive time-dependent receiver operating curve  $AUC(t)$ . Time-dependent concordance is calculated as  $C_{td} = \frac{\int_t AUC(t)w(t)dt}{\int_t w(t)dt}$ , where  $w(t) = f(t) \cdot S(t)$  and  $f(t)$  is defined as the event rate a time  $t$  and  $S(t)$  is the survival probability at time  $t$ , both of which are estimated using Kaplan-Meier estimator Kaplan & Meier (1958).  $AUC(t)$  is the integration of ROC curve under a time bin, where sensitivity and specificity are following the definitions in Heagerty & Zheng (2005). We should note that the time-dependent C-statistics is a more strict metric, especially for non-piecewise survival models, because it upweighs the long time horizon examples which are generally harder to predict for models like traditional CoxPH.

Model	Integrated Brier Score ↓							
	Mort.	Readm.	PH	ATX	CMG	CONS	EDM	PEFF
SwinUNETR <sub>base</sub>	0.071	0.070	0.072	0.071	0.067	0.073	0.074	0.069
SwinUNETR <sub>base/MTL</sub>	0.081	0.078	0.079	0.081	0.078	0.083	0.081	0.080
SwinUNETR <sub>base/visit</sub>	0.073	0.068	0.072	0.070	<b>0.066</b>	0.073	<b>0.074</b>	0.070
SwinUNETR <sub>base/TTE</sub>	<b>0.069</b>	<b>0.067</b>	<b>0.066</b>	<b>0.069</b>	0.067	<b>0.073</b>	0.079	<b>0.069</b>
DenseNet-121 <sub>base</sub>	0.219	0.803	0.702	0.779	0.217	0.533	0.335	0.822
DenseNet-121 <sub>base/MTL</sub>	0.071	0.067	0.068	0.060	0.065	0.072	0.083	0.068
DenseNet-121 <sub>base/visit</sub>	0.063	0.068	0.071	0.071	0.060	0.075	0.075	0.072
DenseNet-121 <sub>base/TTE</sub>	<b>0.028</b>	<b>0.021</b>	<b>0.022</b>	<b>0.014</b>	<b>0.014</b>	<b>0.029</b>	<b>0.023</b>	<b>0.014</b>
ResNet-152 <sub>base</sub>	0.335	0.078	0.479	0.174	0.160	0.166	0.140	0.095
ResNet-152 <sub>base/MTL</sub>	0.076	0.074	0.078	0.075	0.074	0.077	0.078	0.077
ResNet-152 <sub>base/visit</sub>	0.075	0.077	0.074	0.070	0.067	0.077	0.074	0.075
ResNet-152 <sub>base/TTE</sub>	<b>0.068</b>	<b>0.069</b>	<b>0.066</b>	<b>0.065</b>	<b>0.064</b>	<b>0.070</b>	<b>0.070</b>	<b>0.066</b>

Table 5: Prognostic TTE performance for INSPECT, measured by Integrated Brier Score. **Bold** indicates the best performer. Underlined indicates no statistically significant difference versus the \*base/TTE models.

918  
919  
920  
921  
922  
923  
924  
925  
926  
927  
928  
929  
930  
931  
932  
933  
934  
935  
936  
937  
938  
939  
940  
941  
942  
943  
944  
945  
946  
947  
948  
949  
950  
951  
952  
953  
954  
955  
956  
957  
958  
959  
960  
961  
962  
963  
964  
965  
966  
967  
968  
969  
970  
971

Model	Time-dependent C-Statistics $\uparrow$							
	Mort.	Readm.	PH	ATX	CMG	CONS	EDM	PEFF
SwinUNETR <sub>base</sub>	0.649	0.620	0.636	0.602	0.619	0.555	0.637	0.619
SwinUNETR <sub>base/MTL</sub>	0.621	0.624	0.631	0.619	0.625	0.631	0.649	0.633
SwinUNETR <sub>base/visit</sub>	0.625	0.645	0.635	0.650	0.648	0.644	0.648	0.638
SwinUNETR <sub>base/TTE</sub>	<b>0.672</b>	<b>0.657</b>	<b>0.645</b>	<b>0.662</b>	<b>0.660</b>	<b>0.663</b>	<b>0.667</b>	<b>0.667</b>
DenseNet-121 <sub>base</sub>	0.501	0.504	0.542	0.501	0.502	0.508	0.500	0.531
DenseNet-121 <sub>base/MTL</sub>	0.568	0.587	0.585	0.584	0.586	0.586	0.581	0.573
DenseNet-121 <sub>base/visit</sub>	0.632	0.627	0.637	0.624	0.618	0.612	0.624	0.629
DenseNet-121 <sub>base/TTE</sub>	<b>0.677</b>	<b>0.657</b>	<b>0.658</b>	<b>0.651</b>	<b>0.644</b>	<b>0.651</b>	<b>0.658</b>	<b>0.657</b>
ResNet-152 <sub>base</sub>	0.503	0.557	0.517	0.580	0.505	0.563	0.501	0.531
ResNet-152 <sub>base/MTL</sub>	0.632	0.620	0.621	0.615	0.604	0.620	0.618	0.601
ResNet-152 <sub>base/visit</sub>	0.622	0.636	0.621	0.638	0.631	0.635	<u>0.640</u>	0.641
ResNet-152 <sub>base/TTE</sub>	<b>0.678</b>	<b>0.656</b>	<b>0.674</b>	<b>0.667</b>	<b>0.662</b>	<b>0.664</b>	<b>0.670</b>	<b>0.669</b>

Table 6: Prognostic TTE performance for INSPECT, measured by time-dependent C-statistics. **Bold** indicates the best performer. Underlined indicates no statistically significant difference versus the \*<sub>base/TTE</sub> models.

## D MERLIN RESNET-152 RESULTS

For experimental consistency, we selected the ResNet-152<sub>base</sub> model as the starting point for our continued pretraining experiments, however we also conducted an initial study on the Merlin 3D foundation model (Blankemeier et al., 2024). Merlin employs a dual-objective pretraining strategy: first, it uses contrastive loss on radiology notes associated with abdominal CT scans, followed by supervised training where EHR diagnosis codes are used as phenotype classification labels. We compare their reported best performing backbone, ResNet-152, with our weight-inflated (2D ImageNet weights) baseline ResNet-152<sub>base</sub>. Table 7 reports adaption performance for *prognostic classification*, *prognostic TTE* and *diagnostic classification* tasks, where ResNet-152<sub>base</sub> performs similarly to Merlin. While Merlin performs better in mortality classification (though not in the TTE formulation), most tasks show no statistically significant differences across the corresponding metrics. Note that Merlin’s pretraining data consists solely of abdominal CT scans, which likely limits its generalizability to a broader CT data distribution, particularly for capturing lung disease features during pretraining.

Task Category	Task	Metric	Merlin	ResNet <sub>base</sub>	Best
<i>Prognostic Classification</i>	PE (0M)	AUROC	0.518	<b>0.573</b>	ResNet <sub>base</sub>
	Mortality (1M)		<b>0.632</b>	0.583	Merlin
	Mortality (6M)		<b>0.586</b>	0.557	Merlin
	Mortality (12M)		<b>0.590</b>	0.554	Merlin
	Readmission (1M)		<b>0.538</b>	0.536	-
	Readmission (6M)		0.513	<b>0.537</b>	ResNet <sub>base</sub>
	Readmission (12M)		<b>0.538</b>	0.509	-
	PH (12M)		<u>0.514</u>	<b>0.537</b>	-
<i>Prognostic TTE</i>	Mortality	Harrell’s C-index	<b>0.549</b>	0.505	-
	Readmission		0.545	<b>0.560</b>	ResNet <sub>base</sub>
	PH		<b>0.615</b>	0.505	-
	ATX		<u>0.565</u>	<b>0.577</b>	ResNet <sub>base</sub>
	CMG		<b>0.575</b>	0.505	Merlin
	CONS		<b>0.601</b>	0.559	-
	EDM		<b>0.601</b>	0.505	Merlin
	PEFF		<b>0.549</b>	0.536	-
<i>Diagnostic Classification</i>	PE Left	AUROC	0.603	<b>0.654</b>	ResNet <sub>base</sub>
	PE Cent.		0.572	<b>0.687</b>	ResNet <sub>base</sub>
	PE Right		0.562	<b>0.618</b>	ResNet <sub>base</sub>
	PE Chronic		<b>0.550</b>	<u>0.507</u>	-
	PE Acute		<u>0.647</u>	<b>0.695</b>	-
	PE Indeterminate		<u>0.693</u>	<b>0.704</b>	-
	RV/LV <1		0.515	<b>0.593</b>	ResNet <sub>base</sub>
	RV/LV ≥1		0.572	<b>0.626</b>	ResNet <sub>base</sub>

Table 7: Performance comparison of Merlin and ResNet<sub>base</sub> across various tasks. Each task is evaluated using either AUROC for binary classification tasks or Harrell’s C-index for TTE tasks. **Bold** indicates the best performer. Underlined indicates no statistically significant difference between models (i.e.,  $p > 0.05$ ). The Best column denotes the model with the highest performance for each task. Note: Merlin is pretrained on abdominal CTs.

## E MODELING HYPERPARAMETERS

We have detailed our hyperparameter set, both in the image backbones and the probing methods, in Table 8.

Model	Hyperparameters	Values
Image backbones	<b>SwinUNETR</b>	
	learning rate	$10^{-4}, 10^{-5}, 10^{-6}$
	dropout prob	0.1, 0.2, 0.3
	patch size	2x2x2
	window size	7x7x7, 8x8x8
	augmentation strategies	Random rotations, flips
	<b>DenseNet</b>	
	learning rate	$10^{-3}, 10^{-4}, 10^{-5}, 10^{-6}$
	depth	121, 169
	num. of dense blocks	3,4
	dropout prob	0.1, 0.2, 0.3
	augmentation strategies	Random rotations, flips
	<b>ResNet</b>	
	learning rate	$10^{-3}, 10^{-4}, 10^{-5}, 10^{-6}$
	depth	152
residual block	BottleneckBlock	
dropout prob	0.1, 0.2, 0.3	
augmentation strategies	Random rotations, flips	
Probing heads	<b>Logistic Regression</b>	
	penalty	L-1, L-2
	learning rate	0.01, 0.1, 0.2
	solver	lbfgs, liblinear
	scoring	roc_auc, accuracy
	software version	scikit-learn 1.3.2
	<b>DeepSurv</b>	
	num. nodes per layer	32, 64, 128, 256
	depth	2, 3, 4
	learning rate	$10^{-4}, 10^{-5}, 10^{-6}$
dropout prob	0.1, 0.2, 0.3	
software version	pycox 0.2.3	

Table 8: Hyperparameter search grids and software versions of methods under comparison.

## F CLINICAL TASK SUMMARY

Since INSPECT and RSPECT both focus on patients suspected of having pulmonary embolisms, we focus on pulmonary-related pathologies in our experiments. Table 9 outlines the clinical significance of these tasks.

Abbr.	Name	Clinical Meaning
PH	Pulmonary Hypertension	A condition characterized by high blood pressure in the pulmonary arteries, leading to right heart strain and potentially heart failure. Predicting on long term outcome is of higher clinical usage.
PE	Pulmonary Embolism	A life-threatening condition where a blood clot blocks the pulmonary arteries, leading to impaired blood flow to the lungs and acute respiratory symptoms.
ATX	Atelectasis	A collapse of lung tissue that prevents normal oxygen absorption, often resulting from obstruction or fluid accumulation.
CMG	Cardiomegaly	An abnormal enlargement of the heart, typically caused by high blood pressure or heart valve disease, which can lead to heart failure.
CONS	Consolidation	A region of lung tissue filled with liquid instead of air, often seen in pneumonia, causing reduced oxygen exchange.
EDM	Edema	The abnormal accumulation of fluid in tissues or alveolar spaces, often related to heart failure or lung injury, causing difficulty in breathing.
PEFF	Pleural Effusion	The buildup of excess fluid between the layers of the pleura surrounding the lungs, often due to infection, heart failure, or malignancy, leading to difficulty breathing.

Table 9: Clinical definitions of pulmonary conditions used to define TTE tasks in the INSPECT dataset.

## G base PRETRAINING DATASETS

The pretrained image backbone  $_{base}$ , i.e. SwinUNETR, DenseNet and ResNet, have their corresponding pretrained dataset before we take them in for continued pretraining. The details are described in Table 10.

Architecture	Method	Loss	Dim.	Dataset	Size
SwinUNETR $_{base}$	Self-supervised	MAE	3D	Custom Medical	10,050
DenseNet-121 $_{base}$	Supervised	BCE	2D	ImageNet	1,281,167
ResNet-152 $_{base}$	Supervised	BCE	2D	ImageNet	1,281,167

Table 10: Summary of model architectures, pretraining approaches, and off-the-shelf weights. Custom Medical includes 3D Chest, Abdomen, and Head/Neck CT and MRI scans (Valanarasu et al., 2023).

## H EXPERIMENT RUNTIME COST

For our experiments, we have trained the models in a PHI-compliant on-premise server that uses 2 compute nodes, each with 24 Intel Xeon 2.7GHz CPU cores, 200 GB RAM and 4 Nvidia A100 GPUs or 4 H100 GPUs. The total estimate cost for our model  $_{base/TTE}$  regime pretraining is shown in Table 11.

Architecture	Number of GPUs	Estimated wall-clock time	Estimated GPU hours
SwinUNETR $_{base/TTE}$	4 H100 (80GB)	15 days	1,440 GPU hours
DenseNet-121 $_{base/TTE}$	4 A100 (40GB)	9 days	864 GPU hours
ResNet-152 $_{base/TTE}$	4 A100 (80GB)	10 days	960 GPU hours

Table 11: Summary of compute cost across architectures

## I FULL PARAMETER FINE-TUNING ANALYSIS

To illustrate the data efficiency of TTE pretraining method, we conduct single task full parameter fine tuning on both INSPECT and RSPECT dataset, starting from  $\text{base}$  model checkpoint. This experiment is to show what is the cost to take off-the-shelf model to evaluate on the same set of tasks as ours but following traditional machine learning pipeline to conduct full parameter fine-tuning towards each task. Due to the high number of tasks, i.e. 7 binary classification tasks and 8 TTE tasks from INSPECT data, as well as 8 binary classification tasks from RSPECT data, we only choose one representative tasks from each data set and conduct the fine-tuning. We train until the model has early stopping till the plateau on validation set measured by cross entropy loss. The results are shown in Table 12. We can conclude that single task full parameter fine-tuning does not scale as well as TTE pretraining, and in general the performance is no better than, if not much worse than TTE pretraining.

Dataset (fine-tuning task)	Architecture	Full param fine-tuned results (AUROC)	Linear probe results (AUROC)	Delta
INSPECT (12 month PH)	SwinUNETR <sub>base</sub>	0.616	0.597	+ 3.18 %
	SwinUNETR <sub>base/TTE</sub>	0.637	0.672	- 5.20 %
	DenseNet-121 <sub>base</sub>	0.553	0.506	+ 9.28 %
	DenseNet-121 <sub>base/TTE</sub>	0.631	0.689	- 8.41 %
	ResNet-152 <sub>base</sub>	0.552	0.537	+ 2.79 %
	ResNet-152 <sub>base/TTE</sub>	0.649	0.718	- 9.61 %
RSPECT (RV/LV $\geq$ 1)	SwinUNETR <sub>base</sub>	0.606	0.578	+ 4.84 %
	SwinUNETR <sub>base/TTE</sub>	0.631	0.636	- 0.79 %
	DenseNet-121 <sub>base</sub>	0.643	0.607	+ 5.93 %
	DenseNet-121 <sub>base/TTE</sub>	0.611	0.686	- 10.93 %
	ResNet-152 <sub>base</sub>	0.653	0.626	+ 4.31 %
	ResNet-152 <sub>base/TTE</sub>	0.588	0.708	- 16.95 %

Table 12: Summary of compute cost across architectures for full parameter fine-tuning on one single task from each dataset. On average the full parameter fine-tuning computation cost for both INSPECT and RSPECT data ranges from 57.5 (standard deviation: 61.4) to 31.3 (standard deviation: 16.7) GPU hours respectively, which is much more expensive than linear probe, i.e. logistic regression on CPU, which takes minute level computation cost. In addition, the models when fine-tuned on single label task, tend to overfit given the large amount of parameter counts. The above results are from more stringent regularization training.

## J CONFIDENCE INTERVALS

Here we are showing the confidence interval differences between all the baselines against the proposed method ( $_{\text{base}/\text{TTE}}$ ) in both classification tasks (measured by AUROC, in Table 13) and time-to-event tasks (measured by Harrell’s C-index, in Table 14, time-dependent C-statistics, in Table 15 and integrated Brier score in Table 16). The RSPECT dataset confidence interval difference is detailed in Table 17.

Model	Mortality			Readmission			PH
	1M	6M	12M	1M	6M	12M	12M
SwinUNETR <sub>base</sub>	(-0.151, -0.055)*	(-0.156, -0.082)*	(-0.193, -0.113)*	(-0.153, -0.011)*	(-0.093, -0.001)*	(-0.072, -0.011)*	(-0.084, -0.005)*
SwinUNETR <sub>base/MTL</sub>	(-0.103, -0.023)*	(-0.123, -0.053)*	(-0.112, -0.042)*	(-0.140, -0.013)*	(-0.093, -0.002)*	(-0.074, -0.005)*	(-0.083, -0.001)*
SwinUNETR <sub>base/visit</sub>	(-0.123, -0.034)*	(-0.141, -0.064)*	(-0.124, -0.055)*	(-0.135, 0.005)	(-0.092, -0.004)*	(-0.075, -0.004)*	(-0.085, -0.011)*
SwinUNETR <sub>base/TTE</sub>	(0,0)	(0,0)	(0,0)	(0,0)	(0,0)	(0,0)	(0,0)
DenseNet-121 <sub>base</sub>	(-0.193, -0.115)*	(-0.155, -0.092)*	(-0.173, -0.114)*	(-0.185, -0.096)*	(-0.114, -0.044)*	(-0.123, -0.063)*	(-0.161, -0.094)*
DenseNet-121 <sub>base/MTL</sub>	(-0.195, -0.113)*	(-0.153, -0.095)*	(-0.171, -0.113)*	(-0.194, -0.094)*	(-0.112, -0.044)*	(-0.123, -0.065)*	(-0.162, -0.092)*
DenseNet-121 <sub>base/visit</sub>	(-0.155, -0.062)*	(-0.193, -0.133)*	(-0.193, -0.135)*	(-0.195, -0.095)*	(-0.145, -0.065)*	(-0.154, -0.076)*	(-0.163, -0.093)*
DenseNet-121 <sub>base/TTE</sub>	(0,0)	(0,0)	(0,0)	(0,0)	(0,0)	(0,0)	(0,0)
ResNet-152 <sub>base</sub>	(-0.235, -0.143)*	(-0.257, -0.172)*	(-0.222, -0.164)*	(-0.115, -0.006)*	(-0.116, -0.022)*	(-0.102, -0.016)*	(-0.234, -0.152)*
ResNet-152 <sub>base/MTL</sub>	(-0.174, -0.094)*	(-0.155, -0.085)*	(-0.145, -0.083)*	(-0.093, 0.054)	(-0.113, -0.022)*	(-0.112, -0.025)*	(-0.171, -0.094)*
ResNet-152 <sub>base/visit</sub>	(-0.132, -0.056)*	(-0.123, -0.063)*	(-0.113, -0.056)*	(-0.096, -0.062)*	(-0.132, -0.032)*	(-0.117, -0.037)*	(-0.113, -0.042)*
ResNet-152 <sub>base/TTE</sub>	(0,0)	(0,0)	(0,0)	(0,0)	(0,0)	(0,0)	(0,0)

Table 13: 95% confidence interval **differences** for classification performance on INSPECT dataset for proposed method ( $_{\text{base}/\text{TTE}}$ ) and baselines, measure by AUROC. The \* indicates statistical significance under p-value at 0.05 for null hypothesis. (PE: pulmonary embolism; PH: pulmonary hypertension)

Model	Harrell’s C-Index $\uparrow$							
	Mort.	Readm.	PH	ATX	CMG	CONS	EDM	PEFF
SwinUNETR <sub>base</sub>	(-0.381, -0.178)*	(-0.615, -0.126)*	(-0.925, -0.134)*	(-0.292, -0.072)*	(-0.210, -0.121)*	(-0.272, -0.164)*	(-0.455, -0.190)*	(-0.630, -0.264)*
SwinUNETR <sub>base/MTL</sub>	(-0.726, -0.061)*	(-0.404, -0.065)*	(-0.837, -0.138)*	(-0.162, -0.043)*	(-0.603, -0.002)*	(-0.897, -0.072)*	(-0.608, -0.017)*	(-0.086, -0.080)*
SwinUNETR <sub>base/visit</sub>	(-0.104, -0.011)*	(-0.441, 0.083)	(-0.699, -0.077)*	(-0.184, -0.001)*	(-0.333, -0.078)*	(-0.214, -0.127)*	(-0.605, 0.090)	(-0.776, -0.053)*
SwinUNETR <sub>base/TTE</sub>	(0,0)	(0,0)	(0,0)	(0,0)	(0,0)	(0,0)	(0,0)	(0,0)
DenseNet-121 <sub>base</sub>	(-0.564, -0.063)*	(-0.869, -0.067)*	(-0.805, -0.143)*	(-0.235, -0.206)*	(-0.495, -0.298)*	(-0.603, -0.133)*	(-0.672, -0.166)*	(-0.365, -0.053)*
DenseNet-121 <sub>base/MTL</sub>	(-0.155, -0.073)*	(-0.579, 0.001)*	(-0.596, -0.146)*	(-0.663, -0.120)*	(-0.644, -0.047)*	(-0.576, -0.003)*	(-0.965, -0.348)*	(-0.829, -0.141)*
DenseNet-121 <sub>base/visit</sub>	(-0.340, -0.134)*	(-0.642, -0.108)*	(-0.189, 0.044)	(-0.463, -0.053)*	(-0.672, -0.079)*	(-0.420, -0.058)*	(-0.709, -0.071)*	(-0.273, -0.086)*
DenseNet-121 <sub>base/TTE</sub>	(0,0)	(0,0)	(0,0)	(0,0)	(0,0)	(0,0)	(0,0)	(0,0)
ResNet-152 <sub>base</sub>	(-0.332, -0.164)*	(-0.443, -0.257)*	(-0.275, -0.030)*	(-0.615, -0.228)*	(-0.518, -0.016)*	(-0.332, -0.041)*	(-0.289, -0.023)*	(-0.574, -0.165)*
ResNet-152 <sub>base/MTL</sub>	(-0.225, 0.177)	(-0.131, -0.098)*	(-0.440, -0.047)*	(-0.560, -0.284)*	(-0.834, -0.155)*	(0.568, 0.008)*	(-0.672, -0.273)*	(-0.427, -0.041)*
ResNet-152 <sub>base/visit</sub>	(-0.247, -0.166)*	(-0.541, -0.098)*	(-0.178, -0.012)*	(-0.250, -0.019)*	(-0.227, -0.080)*	(-0.186, -0.086)*	(-0.122, -0.056)*	(-0.055, 0.832)
ResNet-152 <sub>base/TTE</sub>	(0,0)	(0,0)	(0,0)	(0,0)	(0,0)	(0,0)	(0,0)	(0,0)

Table 14: 95% confidence interval **differences** for time-to-event performance on INSPECT dataset for proposed method ( $_{\text{base}/\text{TTE}}$ ) and baselines, measured by Harrell’s C-Index. The \* indicates statistical significance under p-value at 0.05 for null hypothesis.

Model	Time-dependent c-statistics $\uparrow$							
	Mort.	Readm.	PH	ATX	CMG	CONS	EDM	PEFF
SwinUNETR <sub>base</sub>	(-0.051, -0.009)*	(-0.050, -0.010)*	(-0.044, -0.016)*	(-0.107, -0.049)*	(-0.074, -0.014)*	(-0.132, -0.072)*	(-0.060, -0.010)*	(-0.090, -0.030)*
SwinUNETR <sub>base/MTL</sub>	(-0.053, -0.007)*	(-0.055, -0.005)*	(-0.053, -0.005)*	(-0.073, -0.013)*	(-0.050, -0.020)*	(-0.062, -0.032)*	(-0.042, -0.012)*	(-0.049, -0.019)*
SwinUNETR <sub>base/visit</sub>	(-0.040, -0.020)*	(-0.037, -0.023)*	(-0.045, -0.015)*	(-0.052, -0.002)*	(-0.049, -0.019)*	(-0.048, -0.018)*	(-0.034, -0.004)*	(-0.033, -0.003)*
SwinUNETR <sub>base/TTE</sub>	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)
DenseNet-121 <sub>base</sub>	(-0.213, -0.143)*	(-0.222, -0.152)*	(-0.146, -0.126)*	(-0.194, -0.144)*	(-0.185, -0.135)*	(-0.163, -0.123)*	(-0.189, -0.149)*	(-0.181, -0.131)*
DenseNet-121 <sub>base/MTL</sub>	(-0.109, -0.069)*	(-0.100, -0.070)*	(-0.093, -0.063)*	(-0.090, -0.060)*	(-0.082, -0.052)*	(-0.095, -0.065)*	(-0.093, -0.063)*	(-0.094, -0.064)*
DenseNet-121 <sub>base/visit</sub>	(-0.075, -0.045)*	(-0.060, -0.030)*	(-0.051, -0.021)*	(-0.063, -0.033)*	(-0.050, -0.020)*	(-0.060, -0.030)*	(-0.064, -0.034)*	(-0.058, -0.028)*
DenseNet-121 <sub>base/TTE</sub>	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)
ResNet-152 <sub>base</sub>	(-0.258, -0.188)*	(-0.099, -0.069)*	(-0.172, -0.142)*	(-0.102, -0.072)*	(-0.247, -0.217)*	(-0.106, -0.076)*	(-0.192, -0.162)*	(-0.188, -0.158)*
ResNet-152 <sub>base/MTL</sub>	(-0.067, -0.037)*	(-0.086, -0.056)*	(-0.063, -0.033)*	(-0.067, -0.037)*	(-0.077, -0.047)*	(-0.059, -0.029)*	(-0.061, -0.031)*	(-0.078, -0.048)*
ResNet-152 <sub>base/visit</sub>	(-0.062, -0.032)*	(-0.035, -0.005)*	(-0.063, -0.033)*	(-0.059, -0.029)*	(-0.046, -0.016)*	(-0.044, -0.014)*	(-0.029, 0.001)	(-0.028, -0.002)*
DenseNet-121 <sub>base/TTE</sub>	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)

Table 15: 95% confidence interval **differences** for time-to-event performance on INSPECT dataset for proposed method ( $_{base/TTE}$ ) and baselines, measured by time-dependent C-statistics. The \* indicates statistical significance under p-value at 0.05 for null hypothesis.

Model	Integrated Brier Score $\downarrow$							
	Mort.	Readm.	PH	ATX	CMG	CONS	EDM	PEFF
SwinUNETR <sub>base</sub>	(0.004, 0.016)*	(0.006, 0.014)*	(0.005, 0.017)*	(0.015, 0.035)*	(0.010, 0.030)*	(0.021, 0.041)*	0.020 (0.010, 0.030)*	(0.012, 0.032)*
SwinUNETR <sub>base/MTL</sub>	(0.002, 0.022)*	(0.001, 0.021)*	(0.003, 0.023)*	(0.003, 0.027)*	(-0.004, 0.026)	(0.005, 0.025)*	(0.013, 0.017)*	(-0.004, 0.026)
SwinUNETR <sub>base/visit</sub>	(-0.006, 0.014)	(0.009, 0.011)*	(0.005, 0.017)*	(0.014, 0.016)*	(0.016, 0.014)*	(0.015, 0.015)*	(-0.020, 0.010)	(0.014, 0.016)*
SwinUNETR <sub>base/TTE</sub>	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)
DenseNet-121 <sub>base</sub>	(0.176, 0.206)*	(0.736, 0.786)*	(0.663, 0.697)*	(0.751, 0.779)*	(0.188, 0.218)*	(0.485, 0.523)*	(0.295, 0.329)*	(0.794, 0.822)*
DenseNet-121 <sub>base/MTL</sub>	(0.033, 0.053)*	(0.031, 0.061)*	(0.029, 0.063)*	(0.039, 0.069)*	(0.036, 0.066)*	(0.024, 0.062)*	(0.041, 0.079)*	(0.039, 0.069)*
DenseNet-121 <sub>base/visit</sub>	(0.032, 0.062)*	(0.032, 0.062)*	(0.032, 0.066)*	(0.042, 0.072)*	(0.039, 0.069)*	(0.027, 0.065)*	(0.033, 0.071)*	(0.043, 0.073)*
DenseNet-121 <sub>base/TTE</sub>	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)
ResNet-152 <sub>base</sub>	(0.266, 0.296)*	(0.006, 0.024)*	(0.412, 0.442)*	(0.094, 0.124)*	(0.086, 0.116)*	(0.081, 0.111)*	(0.055, 0.085)*	(0.028, 0.058)*
ResNet-152 <sub>base/MTL</sub>	(0.001, 0.019)*	(0.005, 0.015)*	(0.003, 0.027)*	(0.005, 0.025)*	(0.005, 0.025)*	(0.008, 0.022)*	(0.007, 0.023)*	(-0.004, 0.026)
ResNet-152 <sub>base/visit</sub>	(0.003, 0.017)*	(0.002, 0.018)*	(0.007, 0.023)*	(0.010, 0.020)*	(0.012, 0.018)*	(0.008, 0.022)*	(0.011, 0.019)*	(0.010, 0.020)*
ResNet-152 <sub>base/TTE</sub>	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)

Table 16: 95% confidence interval **differences** time-to-event performance on INSPECT dataset for proposed method ( $_{base/TTE}$ ) and baselines, measured by integrated Brier score. The \* indicates statistical significance under p-value at 0.05 for null hypothesis.

Model	PE						RV/LV Ratio	
	Left	Cent.	Right	Chronic	Acute	Indet.	<1	$\geq 1$
SwinUNETR <sub>base</sub>	(-0.094, -0.029)*	(-0.139, -0.012)*	(-0.089, -0.035)*	(-0.159, 0.108)	(-0.274, -0.010)*	(-0.131, -0.006)*	(-0.160, -0.075)*	(-0.092, -0.021)*
SwinUNETR <sub>base/MTL</sub>	(-0.097, -0.006)*	(-0.089, 0.058)	(-0.085, -0.005)*	(-0.103, 0.067)	(-0.208, -0.030)*	(-0.158, 0.060)	(-0.174, -0.071)*	(-0.084, -0.009)*
SwinUNETR <sub>base/visit</sub>	(-0.156, -0.017)*	(-0.096, 0.115)	(-0.122, 0.011)	(-0.120, 0.178)	(-0.290, 0.127)	(-0.206, 0.144)	(-0.203, -0.047)*	(-0.210, -0.031)*
SwinUNETR <sub>base/TTE</sub>	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)
DenseNet-121 <sub>base</sub>	(-0.093, -0.012)*	(-0.170, -0.017)*	(-0.077, -0.011)*	(-0.098, 0.052)	(-0.046, 0.073)	(-0.129, 0.054)	(-0.082, 0.002)	(-0.132, -0.031)*
DenseNet-121 <sub>base/MTL</sub>	(-0.049, -0.024)*	(-0.119, -0.042)*	(-0.036, -0.016)*	(-0.146, -0.019)*	(-0.015, 0.018)	(-0.082, -0.010)*	(-0.042, -0.013)*	(-0.068, -0.032)*
DenseNet-121 <sub>base/visit</sub>	(-0.091, -0.013)*	(-0.155, -0.009)*	(-0.071, -0.005)*	(-0.088, 0.054)	(-0.040, 0.079)	(-0.107, 0.085)	(-0.093, -0.007)*	(-0.119, -0.018)*
DenseNet-121 <sub>base/TTE</sub>	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)
ResNet-152 <sub>base</sub>	(-0.039, 0.034)	(-0.061, 0.058)	(-0.059, 0.015)	(-0.032, 0.043)	(-0.026, 0.233)	(-0.116, 0.068)	(-0.048, 0.038)	(-0.142, -0.026)*
ResNet-152 <sub>base/MTL</sub>	(-0.040, 0.037)	(-0.061, 0.066)	(-0.057, 0.015)	(-0.034, 0.041)	(-0.012, 0.213)	(-0.120, 0.076)	(-0.055, 0.036)	(-0.122, -0.012)*
ResNet-152 <sub>base/visit</sub>	(-0.021, -0.005)*	(-0.038, -0.012)*	(-0.024, -0.010)*	(0.030, 0.062)*	(0.012, 0.283)*	(-0.010, 0.006)	(-0.062, -0.001)*	(-0.075, -0.018)*
ResNet-152 <sub>base/TTE</sub>	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)

Table 17: 95% confidence interval **differences** classification performance of different methods on RSPECT dataset for diagnosis labels, measured by AUROC. \* indicates the statistical significance under p value = 0.05



## K TASK HEAD CAPACITY ON RANDOMLY INITIALIZED BACKBONES

We are interested in knowing the different probing method’s utility, with random initialization from different model architectures (i.e. **image backbone provides no predictive power thus only relying on probing methods**), how much does probe methods provide for prediction, and what is delta of each model from that initialization. The comparisons are shown in Figures 5, 6, 7.

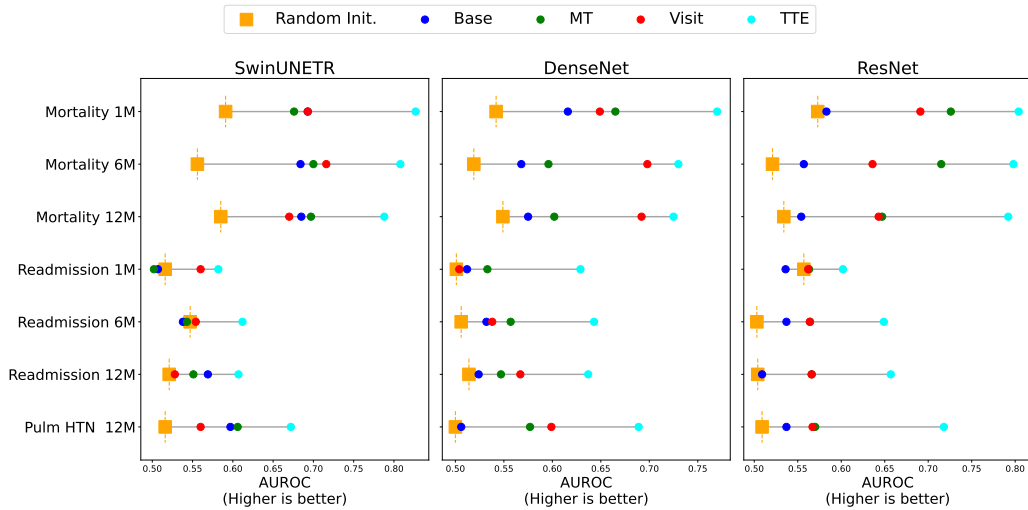


Figure 5: Comparison of models towards random initialization and each model’s delta on AUROC




Figure 6: Comparison of models towards random initialization and each model’s delta on Harrell’s C-index

1350  
 1351  
 1352  
 1353  
 1354  
 1355  
 1356  
 1357  
 1358  
 1359  
 1360  
 1361  
 1362  
 1363  
 1364  
 1365  
 1366  
 1367  
 1368  
 1369  
 1370  
 1371  
 1372  
 1373  
 1374  
 1375  
 1376  
 1377  
 1378  
 1379  
 1380  
 1381  
 1382  
 1383  
 1384  
 1385  
 1386  
 1387  
 1388  
 1389  
 1390  
 1391  
 1392  
 1393  
 1394  
 1395  
 1396  
 1397  
 1398  
 1399  
 1400  
 1401  
 1402  
 1403

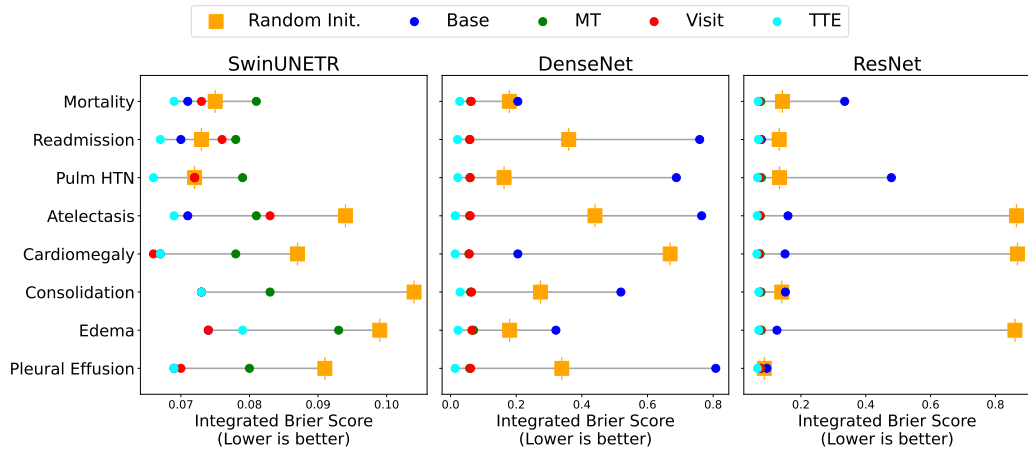


Figure 7: Comparison of models towards random initialization and each model's delta on integrated Brier score

## L PIECEWISE EXPONENTIAL LOSS FUNCTION

In our methods section, we briefly describe that we use a piecewise exponential model, and thus, we use the corresponding piecewise exponential survival loss as our loss function. In this appendix section, we provide the exact formulas used to implement that survival loss.

For the piecewise exponential, we select a number of pieces and each piece covers the response period starting at  $S_p$  and ending at  $E_p$ . For every piece  $p$  and image  $i$  in our dataset, our model needs to output  $\lambda_{ip}$ , the instantaneous hazard rate for the patient described by that image during the response time period  $p$ . These  $\lambda$  values are then used to define a survival function for each image  $i$ :

$$S_i(t) = \prod_{p=1}^P \exp[-\lambda_{ip}(\min(t, E_p) - S_p)]I(t \geq S_p), \quad (3)$$

This survival function has an associated PDF,  $f$ , that is simply the negative of the derivative of the survival function.

$$f_i = -\frac{\partial S_i}{\partial t} \quad (4)$$

Applying the derivative operator, we obtain that

$$f_i = S_i(t) \sum_{p=1}^P \lambda_{ip} I(S_e \leq t \leq S_p) \quad (5)$$

We define the loss function for training by using the survival function to implement the standard survival likelihood equation and taking the product of that likelihood over the entire dataset.  $\Delta_i$  is event indicator whose value is 1 when the actual event is observed. We can then plug in our definitions for the survival function  $S$  and the pdf  $f$  to obtain our final loss function:

$$\begin{aligned} \mathcal{L} &= \prod_{i=1}^n [S_i(t)]^{1-\Delta_i} [f_i(t)]^{\Delta_i} \\ \mathcal{L} &= \prod_{i=1}^n [S_i(t)] \left[ \sum_{p=1}^P \lambda_{ip} I(S_e \leq t \leq S_p) \right]^{\Delta_i} \\ \mathcal{L} &= \prod_{i=1}^n \left[ \prod_{p=1}^P \exp[-\lambda_{ip}(\min(t, E_p) - S_p)] I(t \geq S_p) \right] \left[ \sum_{p=1}^P \lambda_{ip} I(S_e \leq t \leq S_p) \right]^{\Delta_i} \end{aligned} \quad (6)$$

## M KAPLAN-MEIER CURVES FOR TTE TASKS

We here plot the stratified groups in terms of the diagnosis of pulmonary embolism for Kaplan-Meier curves among all of our TTE tasks, shown in Figures 8, 9, 10, 11, 12, 13, 14, 15.




Figure 8: Kaplan-Meier curve for Mortality




Figure 9: Kaplan-Meier curve for Readmission

1512  
1513  
1514  
1515  
1516  
1517  
1518  
1519  
1520  
1521  
1522  
1523  
1524  
1525  
1526  
1527  
1528  
1529  
1530  
1531  
1532  
1533  
1534  
1535  
1536  
1537  
1538  
1539  
1540  
1541  
1542  
1543  
1544  
1545  
1546  
1547  
1548  
1549  
1550  
1551  
1552  
1553  
1554  
1555  
1556  
1557  
1558  
1559  
1560  
1561  
1562  
1563  
1564  
1565




Figure 10: Kaplan-Meier curve for Pulmonary Hypertension (PH)




Figure 11: Kaplan-Meier curve for Atelectasis

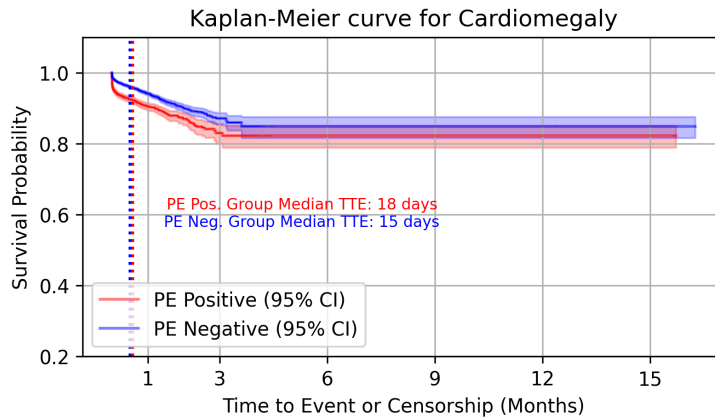


Figure 12: Kaplan-Meier curve for Cardiomegaly

1566  
1567  
1568  
1569  
1570  
1571  
1572  
1573  
1574  
1575  
1576  
1577  
1578  
1579  
1580  
1581  
1582  
1583  
1584  
1585  
1586  
1587  
1588  
1589  
1590  
1591  
1592  
1593  
1594  
1595  
1596  
1597  
1598  
1599  
1600  
1601  
1602  
1603  
1604  
1605  
1606  
1607  
1608  
1609  
1610  
1611  
1612  
1613  
1614  
1615  
1616  
1617  
1618  
1619




Figure 13: Kaplan-Meier curve for Consolidation

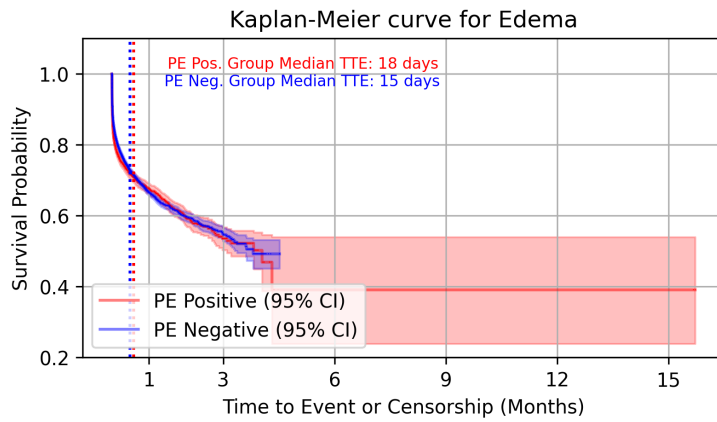


Figure 14: Kaplan-Meier curve for Edema




Figure 15: Kaplan-Meier curve for Pleural Effusion

## N PRETRAINING TASK EVENT TIME DISTRIBUTIONS

We have employed 8,192 future clinical events in our pretraining mixture and here we plot the distribution of the time to event w.r.t to each of the 8,192 events across all the cohorts in INSPECT data. We have maximum and median time to event distributions for the 8192 labels and all time to event distribution for each occurrence in Figures 16, 17, 18.




Figure 16: Maximum Time to Event across INSPECT cohort for 8192 future events (Per event)




Figure 17: Median Time to Event across INSPECT cohort for 8192 future events (Per event)




Figure 18: All Time to Event across INSPECT cohort for 8192 future events (Per occurrence)

## O EXAMPLES OF PRETRAINING CLINICAL EVENTS

We have list the set of medical ontologies that are available in the EHR modality, in Table 18. We further stratify the frequency of our 8,192 pretraining clinical events across all cohorts in the INSPECT dataset into quintiles and present the top 3 examples in each quintile, in Table 19.

Ontology
OMOP Extension
Medicare Specialty
CPT4
CVX
ICD9Proc
RxNorm
SNOMED
RxNorm Extension
Cancer Modifier
ICD10PCS
CMS Place of Service
Visit
Ethnicity
Gender
ICDO3
Race
LOINC
HCPCS

Table 18: OHDSI Athena ontologies used in our benchmark


Medical Code	Description	Quintile
SNOMED/60853003	Disorder of magnesium metabolism	1
SNOMED/167321001	Urine dipstick for urobilinogen	1
LOINC/2472-9	Immunoglobulin M (IgM) in serum or plasma	1
RxNorm/905216	Hydralazine Hydrochloride 50 MG	2
SNOMED/276319003	Finding of menstrual bleeding	2
RxNorm/979463	Losartan potassium 100 MG	2
SNOMED/89659001	Amylase measurement, serum	3
SNOMED/400166009	Acquired keratoderma	3
SNOMED/70209001	Late effect of complications of procedure (disorder)	3
LOINC/32355-0	Bacteria identified in Specimen by Respiratory culture	4
SNOMED/56675007	Acute heart failure	4
SNOMED/297971001	Finding of sensation of skin	4
SNOMED/118664000	Procedure on body system	5
SNOMED/118717007	Procedure on organ	5
SNOMED/118672003	Procedure on cardiovascular system	5

Table 19: Top 3 events in each quintile




**P FEATURE ATTRIBUTION**


We use GradCAM (Selvaraju et al., 2017) to visualize the Gradient-weighted Class Activation Mapping for our DenseNet-121 model and its baselines, w.r.t. to the TTE tasks we curated, where we selected the CTs with the corresponding TTE labels that actually happened for the patient after the CT capture, in Figure 19. We can observe that the TTE version of the model can focus the pathology in a more reasonable region rather than scattered features learned by different baselines.




(a) GradCAM for Atelectasis=True after CT's capture for models and baselines




(b) GradCAM for Consolidation=True after CT's capture for models and baselines




(c) GradCAM for Pleural Effusion=True after CT's capture for models and baselines



(d) GradCAM for Cardiomegaly=True after CT's capture for models and baselines



(e) GradCAM for Edema=True after CT's capture for models and baselines



(f) GradCAM for Pulmonary Hypertension=True after CT's capture for models and baselines

Figure 19: GradCAM visualizations of sampled CTs for various occurred conditions: Atelectasis, Consolidation, Pleural Effusion, Cardiomegaly, Edema, and Pulmonary Hypertension.

## Q SUBGROUP PERFORMANCE

We conduct an analysis to assess performance on our best performing model, SwinUNETR, across sensitive subgroups, ensuring our pretraining technique does not unfairly disadvantage any group compared to existing methods. [The stratified subgroup counts for this experiment under test split of INSPECT dataset is in Table 20.](#) We evaluate performance using the AUROC (bootstrapped,  $n=1000$ ) on 5 binary prognostic tasks, comparing TTE against `base`, showing that TTE pretraining does not reduce performance for sensitive groups and generally improves risk ranking across all groups. The details are in Tables [21](#), [22](#), [23](#), [24](#), [25](#), [26](#), [27](#).

Group	Concept	Counts
Gender	Female	1844
	Male	1370
Age	18-39	448
	39-69	1667
	69-89	993
	>89	106
Race	Asian	554
	Black	192
	Native	74
	White	1730
	Unknown	664
Ethnicity	Hispanic	495
	Not Hispanic	2612
	Unknown	107

Table 20: [Subgroup stratified counts for the test split under INSPECT dataset](#)

Group	Concept	AUROC $\uparrow$			
		SwinUNETR <sub>base</sub>	SwinUNETR <sub>base/MTL</sub>	SwinUNETR <sub>base/visit</sub>	SwinUNETR <sub>base/TTE</sub>
Gender	Female	0.666 (0.607, 0.722)	0.648 (0.588, 0.705)	0.690 (0.636, 0.739)	<b>0.828 (0.789, 0.866)</b>
	Male	0.735 (0.670, 0.800)	0.712 (0.637, 0.785)	0.696 (0.634, 0.753)	<b>0.841 (0.797, 0.879)</b>
Age	18-39	0.753 (0.591, 0.887)	0.715 (0.553, 0.863)	0.643 (0.512, 0.765)	<b>0.887 (0.808, 0.952)</b>
	39-69	0.701 (0.639, 0.755)	0.688 (0.626, 0.751)	0.683 (0.628, 0.735)	<b>0.846 (0.809, 0.882)</b>
	69-89	0.661 (0.580, 0.738)	0.643 (0.558, 0.720)	0.730 (0.672, 0.790)	<b>0.809 (0.752, 0.864)</b>
	>89	0.662 (0.435, 0.869)	0.648 (0.434, 0.877)	<b>0.723 (0.445, 0.916)</b>	0.647 (0.426, 0.847)
Race	Asian	0.675 (0.547, 0.787)	0.642 (0.522, 0.760)	0.668 (0.562, 0.774)	<b>0.831 (0.732, 0.912)</b>
	Black	0.533 (0.328, 0.738)	0.512 (0.313, 0.707)	0.660 (0.523, 0.785)	<b>0.805 (0.690, 0.905)</b>
	Native	0.667 (0.040, 1.000)	0.689 (0.040, 1.000)	0.743 (0.333, 0.997)	<b>0.773 (0.555, 0.939)</b>
	White	0.749 (0.694, 0.797)	0.732 (0.673, 0.786)	0.718 (0.665, 0.769)	<b>0.855 (0.816, 0.889)</b>
	Unknown	0.649 (0.561, 0.733)	0.636 (0.553, 0.726)	0.670 (0.598, 0.740)	<b>0.778 (0.719, 0.830)</b>
Ethnicity	Hispanic	0.645 (0.532, 0.754)	0.646 (0.524, 0.758)	0.685 (0.588, 0.778)	<b>0.785 (0.701, 0.866)</b>
	Not Hispanic	0.701 (0.648, 0.751)	0.681 (0.630, 0.730)	0.699 (0.655, 0.741)	<b>0.842 (0.809, 0.873)</b>
	Unknown	0.731 (0.563, 0.862)	0.720 (0.566, 0.848)	0.688 (0.572, 0.800)	<b>0.832 (0.744, 0.912)</b>

Table 21: 1-Month Mortality prognosis breakdown by subgroups, reported as the mean AUROC of a test set bootstrap ( $n=1000$ ) with 95% CI. **Bold** indicates the best performance across all models.

1836

1837

1838

1839

1840

1841

1842

1843

1844

1845

1846

1847

1848

1849

1850

1851

1852

1853

1854

1855

1856

1857

1858

1859

1860

1861

1862

1863

1864

1865

1866

1867

1868

1869

1870

1871

1872

1873

1874

1875

1876

1877

1878

1879

1880

1881

1882

1883

1884

1885

1886

1887

1888

1889

		AUROC $\uparrow$			
Group	Concept	SwinUNETR <sub>base</sub>	SwinUNETR <sub>base/MTL</sub>	SwinUNETR <sub>base/visit</sub>	SwinUNETR <sub>base/TTE</sub>
Gender	Female	0.661 (0.620, 0.702)	0.676 (0.632, 0.717)	0.699 (0.653, 0.743)	<b>0.813 (0.779, 0.842)</b>
	Male	0.710 (0.657, 0.762)	0.739 (0.688, 0.790)	0.727 (0.683, 0.769)	<b>0.819 (0.779, 0.858)</b>
Age	18-39	0.799 (0.727, 0.861)	0.803 (0.729, 0.867)	0.723 (0.634, 0.810)	<b>0.887 (0.820, 0.944)</b>
	39-69	0.685 (0.641, 0.725)	0.700 (0.653, 0.747)	0.690 (0.645, 0.734)	<b>0.811 (0.775, 0.842)</b>
	69-89	0.635 (0.571, 0.691)	0.657 (0.596, 0.717)	0.755 (0.693, 0.805)	<b>0.799 (0.752, 0.844)</b>
	>89	0.680 (0.502, 0.844)	<b>0.738 (0.562, 0.880)</b>	0.687 (0.473, 0.870)	0.732 (0.528, 0.883)
Race	Asian	0.666 (0.583, 0.740)	0.693 (0.608, 0.771)	0.682 (0.591, 0.770)	<b>0.755 (0.667, 0.840)</b>
	Black	0.573 (0.414, 0.728)	0.615 (0.478, 0.758)	0.666 (0.536, 0.793)	<b>0.818 (0.720, 0.905)</b>
	Native	0.806 (0.500, 1.000)	0.794 (0.488, 0.986)	0.795 (0.586, 0.966)	<b>0.882 (0.710, 0.990)</b>
	White	0.713 (0.674, 0.751)	0.724 (0.684, 0.765)	0.733 (0.694, 0.772)	<b>0.834 (0.802, 0.862)</b>
	Unknown	0.664 (0.612, 0.713)	0.680 (0.627, 0.728)	0.707 (0.650, 0.761)	<b>0.779 (0.739, 0.818)</b>
Ethnicity	Hispanic	0.651 (0.554, 0.743)	0.673 (0.579, 0.762)	0.679 (0.597, 0.756)	<b>0.767 (0.694, 0.833)</b>
	Not Hispanic	0.688 (0.653, 0.724)	0.709 (0.675, 0.743)	0.724 (0.688, 0.759)	<b>0.820 (0.795, 0.845)</b>
	Unknown	0.693 (0.624, 0.758)	0.689 (0.625, 0.753)	0.726 (0.647, 0.795)	<b>0.792 (0.741, 0.840)</b>

Table 22: 6-Month Mortality prognosis breakdown by subgroups, reported as the mean AUROC of a test set bootstrap (n=1000) with 95% CI. **Bold** indicates the best performance across all models.

1854

1855

1856

1857

1858

1859

1860

1861

1862

1863

1864

1865

1866

1867

1868

1869

1870

1871

1872

1873

1874

1875

1876

1877

1878

1879

1880

1881

1882

1883

1884

1885

1886

1887

1888

1889

		AUROC $\uparrow$			
Group	Concept	SwinUNETR <sub>base</sub>	SwinUNETR <sub>base/MTL</sub>	SwinUNETR <sub>base/visit</sub>	SwinUNETR <sub>base/TTE</sub>
Gender	Female	0.655 (0.616, 0.694)	0.669 (0.630, 0.706)	0.660 (0.617, 0.702)	<b>0.796 (0.762, 0.825)</b>
	Male	0.717 (0.669, 0.762)	0.732 (0.685, 0.775)	0.672 (0.626, 0.717)	<b>0.796 (0.761, 0.833)</b>
Age	18-39	0.828 (0.760, 0.886)	0.839 (0.777, 0.893)	0.681 (0.589, 0.764)	<b>0.876 (0.805, 0.936)</b>
	39-69	0.682 (0.636, 0.722)	0.694 (0.653, 0.731)	0.644 (0.598, 0.686)	<b>0.791 (0.759, 0.825)</b>
	69-89	0.632 (0.573, 0.687)	0.647 (0.589, 0.702)	0.706 (0.649, 0.763)	<b>0.784 (0.738, 0.827)</b>
	>89	0.605 (0.434, 0.771)	0.616 (0.429, 0.796)	0.632 (0.440, 0.809)	<b>0.660 (0.481, 0.819)</b>
Race	Asian	0.675 (0.604, 0.746)	0.695 (0.624, 0.763)	0.657 (0.567, 0.741)	<b>0.754 (0.684, 0.824)</b>
	Black	0.624 (0.498, 0.746)	0.645 (0.522, 0.766)	0.620 (0.505, 0.730)	<b>0.781 (0.683, 0.866)</b>
	Native	0.716 (0.451, 0.950)	0.690 (0.394, 0.937)	0.745 (0.487, 0.952)	<b>0.840 (0.650, 0.981)</b>
	White	0.716 (0.675, 0.755)	0.724 (0.687, 0.762)	0.676 (0.632, 0.714)	<b>0.822 (0.793, 0.851)</b>
	Unknown	0.657 (0.609, 0.703)	0.673 (0.625, 0.719)	0.672 (0.621, 0.725)	<b>0.750 (0.707, 0.788)</b>
Ethnicity	Hispanic	0.643 (0.553, 0.728)	0.660 (0.571, 0.746)	0.631 (0.551, 0.701)	<b>0.708 (0.630, 0.780)</b>
	Not Hispanic	0.689 (0.657, 0.721)	0.702 (0.668, 0.736)	0.675 (0.640, 0.714)	<b>0.808 (0.782, 0.835)</b>
	Unknown	0.698 (0.638, 0.755)	0.706 (0.647, 0.761)	0.695 (0.621, 0.764)	<b>0.764 (0.713, 0.811)</b>

Table 23: 12-Month Mortality prognosis breakdown by subgroups, reported as the mean AUROC of a test set bootstrap (n=1000) with 95% CI. **Bold** indicates the best performance across all models.

1871

1872

1873

1874

1875

1876

1877

1878

1879

1880

1881

1882

1883

1884

1885

1886

1887

1888

1889

		AUROC $\uparrow$			
Group	Concept	SwinUNETR <sub>base</sub>	SwinUNETR <sub>base/MTL</sub>	SwinUNETR <sub>base/visit</sub>	SwinUNETR <sub>base/TTE</sub>
Gender	Female	0.539 (0.459, 0.623)	0.534 (0.457, 0.614)	0.551 (0.477, 0.630)	<b>0.599 (0.522, 0.671)</b>
	Male	0.498 (0.391, 0.607)	0.492 (0.386, 0.600)	0.532 (0.448, 0.613)	<b>0.587 (0.499, 0.675)</b>
Age	18-39	0.452 (0.280, 0.634)	0.438 (0.260, 0.613)	0.598 (0.456, 0.737)	<b>0.603 (0.423, 0.760)</b>
	39-69	0.458 (0.373, 0.546)	0.457 (0.374, 0.546)	0.523 (0.446, 0.597)	<b>0.558 (0.482, 0.628)</b>
	69-89	0.687 (0.573, 0.794)	<b>0.689 (0.581, 0.787)</b>	0.492 (0.370, 0.608)	0.676 (0.547, 0.787)
	>89	0.459 (0.013, 0.907)	0.432 (0.013, 0.878)	<b>0.772 (0.585, 0.915)</b>	0.363 (0.200, 0.533)
Race	Asian	0.594 (0.442, 0.742)	0.603 (0.428, 0.748)	<b>0.674 (0.485, 0.854)</b>	0.619 (0.460, 0.763)
	Black	0.374 (0.025, 0.848)	0.397 (0.027, 0.865)	<b>0.523 (0.338, 0.775)</b>	0.524 (0.201, 0.945)
	Native	0.860 (0.760, 0.940)	0.880 (0.780, 0.960)	0.287 (0.146, 0.447)	<b>0.941 (0.875, 1.000)</b>
	White	0.490 (0.411, 0.573)	0.478 (0.395, 0.558)	0.521 (0.449, 0.591)	<b>0.614 (0.540, 0.686)</b>
	Unknown	0.491 (0.408, 0.571)	0.489 (0.410, 0.575)	<b>0.613 (0.525, 0.700)</b>	0.525 (0.444, 0.604)
Ethnicity	Hispanic	<b>0.670 (0.479, 0.836)</b>	0.662 (0.468, 0.842)	0.515 (0.406, 0.626)	0.587 (0.400, 0.774)
	Not Hispanic	0.515 (0.439, 0.586)	0.513 (0.442, 0.587)	0.542 (0.475, 0.610)	<b>0.620 (0.557, 0.680)</b>
	Unknown	0.433 (0.347, 0.522)	0.426 (0.340, 0.511)	<b>0.666 (0.563, 0.761)</b>	0.480 (0.380, 0.576)

Table 24: 1-Month Readmission prognosis breakdown by subgroups, reported as the mean AUROC of a test set bootstrap (n=1000) with 95% CI. **Bold** indicates the best performance across all models.

		AUROC $\uparrow$			
Group	Concept	SwinUNETR <sub>base</sub>	SwinUNETR <sub>base/MTL</sub>	SwinUNETR <sub>base/visit</sub>	SwinUNETR <sub>base/TTE</sub>
Gender	Female	0.571 (0.522, 0.619)	0.577 (0.528, 0.626)	0.488 (0.436, 0.541)	<b>0.607 (0.558, 0.654)</b>
	Male	0.525 (0.466, 0.585)	0.526 (0.470, 0.583)	<b>0.608 (0.541, 0.670)</b>	0.607 (0.556, 0.660)
Age	18-39	0.533 (0.428, 0.630)	0.542 (0.440, 0.636)	0.541 (0.434, 0.654)	<b>0.660 (0.569, 0.753)</b>
	39-69	0.516 (0.467, 0.564)	0.519 (0.472, 0.566)	0.519 (0.461, 0.573)	<b>0.590 (0.544, 0.634)</b>
	69-89	0.616 (0.542, 0.691)	<b>0.624 (0.549, 0.694)</b>	0.537 (0.456, 0.616)	0.610 (0.535, 0.689)
	>89	0.604 (0.000, 0.984)	0.616 (0.000, 0.984)	<b>0.792 (0.601, 0.956)</b>	0.553 (0.161, 0.921)
Race	Asian	0.548 (0.456, 0.641)	0.556 (0.459, 0.648)	0.549 (0.401, 0.693)	<b>0.591 (0.497, 0.678)</b>
	Black	0.510 (0.305, 0.706)	0.509 (0.305, 0.698)	0.525 (0.365, 0.697)	<b>0.557 (0.382, 0.734)</b>
	Native	0.841 (0.612, 0.989)	<b>0.850 (0.659, 0.985)</b>	0.598 (0.431, 0.774)	0.835 (0.588, 0.977)
	White	0.554 (0.504, 0.604)	0.561 (0.511, 0.610)	0.529 (0.479, 0.580)	<b>0.626 (0.575, 0.675)</b>
	Unknown	0.504 (0.451, 0.559)	0.502 (0.451, 0.553)	0.591 (0.522, 0.652)	<b>0.594 (0.544, 0.642)</b>
Ethnicity	Hispanic	<b>0.610 (0.516, 0.704)</b>	0.606 (0.510, 0.695)	0.594 (0.503, 0.685)	0.556 (0.455, 0.648)
	Not Hispanic	0.553 (0.507, 0.596)	0.560 (0.519, 0.601)	0.528 (0.475, 0.574)	<b>0.620 (0.581, 0.661)</b>
	Unknown	0.469 (0.401, 0.532)	0.472 (0.413, 0.534)	0.600 (0.511, 0.687)	<b>0.605 (0.541, 0.665)</b>

Table 25: 6-Month Readmission prognosis breakdown by subgroups, reported as the mean AUROC of a test set bootstrap (n=1000) with 95% CI. **Bold** indicates the best performance across all models.

		AUROC $\uparrow$			
Group	Concept	SwinUNETR <sub>base</sub>	SwinUNETR <sub>base/MTL</sub>	SwinUNETR <sub>base/visit</sub>	SwinUNETR <sub>base/TTE</sub>
Gender	Female	0.588 (0.539, 0.629)	0.565 (0.520, 0.609)	0.492 (0.444, 0.538)	<b>0.607 (0.562, 0.650)</b>
	Male	0.563 (0.512, 0.612)	0.547 (0.497, 0.598)	0.555 (0.501, 0.610)	<b>0.601 (0.552, 0.650)</b>
Age	18-39	0.557 (0.467, 0.645)	0.554 (0.462, 0.639)	0.496 (0.402, 0.581)	<b>0.653 (0.573, 0.730)</b>
	39-69	0.562 (0.519, 0.605)	0.530 (0.484, 0.571)	0.509 (0.461, 0.557)	<b>0.596 (0.554, 0.640)</b>
	69-89	<b>0.611 (0.544, 0.673)</b>	0.600 (0.529, 0.664)	0.529 (0.464, 0.596)	0.592 (0.523, 0.659)
	>89	0.589 (0.019, 0.913)	0.675 (0.180, 0.981)	<b>0.678 (0.466, 0.875)</b>	0.564 (0.130, 0.855)
Race	Asian	0.614 (0.524, 0.694)	0.587 (0.497, 0.679)	0.473 (0.383, 0.561)	<b>0.635 (0.555, 0.716)</b>
	Black	0.494 (0.336, 0.680)	0.486 (0.325, 0.642)	0.488 (0.344, 0.635)	<b>0.634 (0.472, 0.787)</b>
	Native	<b>0.901 (0.763, 0.987)</b>	0.843 (0.679, 0.976)	0.631 (0.462, 0.792)	0.861 (0.658, 1.000)
	White	0.589 (0.542, 0.632)	0.564 (0.519, 0.607)	0.516 (0.472, 0.562)	<b>0.626 (0.585, 0.668)</b>
	Unknown	0.523 (0.481, 0.569)	0.515 (0.470, 0.560)	<b>0.566 (0.510, 0.621)</b>	0.560 (0.512, 0.607)
Ethnicity	Hispanic	0.522 (0.440, 0.604)	0.534 (0.447, 0.622)	<b>0.602 (0.510, 0.689)</b>	0.479 (0.401, 0.557)
	Not Hispanic	0.597 (0.560, 0.634)	0.570 (0.534, 0.608)	0.506 (0.464, 0.545)	<b>0.635 (0.601, 0.669)</b>
	Unknown	0.509 (0.454, 0.565)	0.501 (0.437, 0.557)	<b>0.554 (0.472, 0.631)</b>	0.589 (0.529, 0.644)

Table 26: 12-Month Readmission prognosis breakdown by subgroups, reported as the mean AUROC of a test set bootstrap (n=1000) with 95% CI. **Bold** indicates the best performance across all models.

		AUROC $\uparrow$			
Group	Concept	SwinUNETR <sub>base</sub>	SwinUNETR <sub>base/MTL</sub>	SwinUNETR <sub>base/visit</sub>	SwinUNETR <sub>base/TTE</sub>
Gender	Female	0.597 (0.547, 0.645)	0.607 (0.556, 0.655)	0.559 (0.515, 0.602)	<b>0.665 (0.618, 0.710)</b>
	Male	0.563 (0.504, 0.622)	0.572 (0.512, 0.631)	0.542 (0.496, 0.592)	<b>0.673 (0.625, 0.721)</b>
Age	18-39	0.543 (0.379, 0.709)	0.553 (0.386, 0.703)	0.515 (0.431, 0.594)	<b>0.758 (0.650, 0.862)</b>
	39-69	0.606 (0.552, 0.660)	0.615 (0.562, 0.669)	0.552 (0.506, 0.596)	<b>0.661 (0.614, 0.705)</b>
	69-89	0.547 (0.486, 0.609)	0.556 (0.497, 0.616)	0.579 (0.516, 0.639)	<b>0.644 (0.584, 0.705)</b>
	>89	0.539 (0.361, 0.731)	<b>0.548 (0.361, 0.728)</b>	0.534 (0.404, 0.669)	0.464 (0.294, 0.631)
Race	Asian	0.613 (0.521, 0.707)	0.620 (0.523, 0.714)	0.547 (0.470, 0.627)	<b>0.667 (0.576, 0.750)</b>
	Black	0.706 (0.532, 0.850)	0.727 (0.562, 0.876)	0.373 (0.249, 0.501)	<b>0.767 (0.644, 0.872)</b>
	Native	0.378 (0.074, 0.711)	0.326 (0.046, 0.644)	0.218 (0.075, 0.369)	<b>0.583 (0.312, 0.833)</b>
	White	0.579 (0.529, 0.630)	0.593 (0.543, 0.643)	0.580 (0.539, 0.625)	<b>0.687 (0.643, 0.730)</b>
	Unknown	0.603 (0.558, 0.647)	0.609 (0.566, 0.655)	0.581 (0.532, 0.633)	<b>0.641 (0.595, 0.686)</b>
Ethnicity	Hispanic	0.602 (0.511, 0.687)	0.615 (0.521, 0.703)	0.500 (0.414, 0.584)	<b>0.641 (0.550, 0.722)</b>
	Not Hispanic	0.584 (0.538, 0.626)	0.589 (0.547, 0.632)	0.562 (0.526, 0.598)	<b>0.679 (0.643, 0.715)</b>
	Unknown	0.618 (0.566, 0.672)	0.626 (0.570, 0.677)	0.597 (0.526, 0.662)	<b>0.646 (0.594, 0.696)</b>

Table 27: 12-Month Pulmonary Hypertension prognosis breakdown by subgroups, reported as the mean AUROC of a test set bootstrap (n=1000) with 95% CI. **Bold** indicates the best performance across all models.