
CORE-BEHRT

A CAREFULLY OPTIMIZED AND RIGOROUSLY EVALUATED BEHRT

A PREPRINT

Mikkel Odgaard^{*†1}, Kiril Vadimovic Klein^{*†1}, Sanne Møller Thysen², Espen Jimenez-Solem^{2,3}, Martin Sillesen⁴, and Mads Nielsen¹

¹Department of Computer Science, University of Copenhagen

²Department of Clinical Pharmacology, Copenhagen University Hospital Bispebjerg and Frederiksberg

³Department of Clinical Medicine, University of Copenhagen

⁴Department of Organ Surgery and Transplantation, Copenhagen University Hospital, Rigshospitalet
Copenhagen, Denmark

^{*}Corresponding Authors: miod.di.ku.dk, kikl@di.ku.dk

[†]Equal Contribution

ABSTRACT

BERT-based models for Electronic Health Records (EHR) have surged in popularity following the release of BEHRT and Med-BERT. Subsequent models have largely built on these foundations despite the fundamental design choices of these pioneering models remaining underexplored. To address this issue, we introduce CORE-BEHRT, a Carefully Optimized and Rigorously Evaluated BEHRT. Through incremental optimization, we isolate the sources of improvement for key design choices, giving us insights into the effect of data representation and individual technical components on performance. Evaluating this across a set of generic tasks (death, pain treatment, and general infection), we showed that improving data representation can increase the average downstream performance from 0.785 to 0.797 AUROC ($p < 10^{-7}$), primarily when including medication and timestamps. Improving the architecture and training protocol on top of this increased average downstream performance to 0.801 AUROC ($p < 10^{-7}$). We then demonstrated the consistency of our optimization through a rigorous evaluation across 25 diverse clinical prediction tasks. We observed significant performance increases in 17 out of 25 tasks and improvements in 24 tasks, highlighting the generalizability of our findings. Our findings provide a strong foundation for future work and aim to increase the trustworthiness of BERT-based EHR models.

Keywords BEHRT, Med-BERT, BERT, Transformers, EHR, Risk Prediction, Deep Learning, Real World Data

1 Introduction

In recent years, leveraging Electronic Health Records (EHR) to predict patient outcomes has gained significant interest. Initially, progress was gradual due to a select number of available machine learning methods, such as logistic regression and random forest [MJ et al., 1997, Khalilia et al., 2011]. However, the advent of deep learning, coupled with a drastic increase in available EHR data [Slawomirski et al., 2023], marked a pivotal shift and enabled the extraction of more complex patterns and insights.

A notable development in utilizing EHR data more effectively has been the adoption of architectures originally developed for natural language processing (NLP) and computer vision. This has been exemplified by the use of recurrent [Choi et al., 2015, 2016a,b] and convolutional [Nguyen et al., 2016] neural networks on EHR data. The introduction of the transformer architecture [Vaswani et al., 2017], known for its versatility across different data types, significantly impacted EHR analysis. Specifically, building on the seminal work of BERT (Bidirectional Encoder Representations from Transformers) [Devlin et al., 2018], the field has seen the emergence of transformer-based models tailored to EHR data [Li et al., 2020, Rasmy et al., 2021]. These early models inspired a wave of subsequent research introducing a variety of modifications touching on input data representation, model architecture, downstream tasks, and evaluation

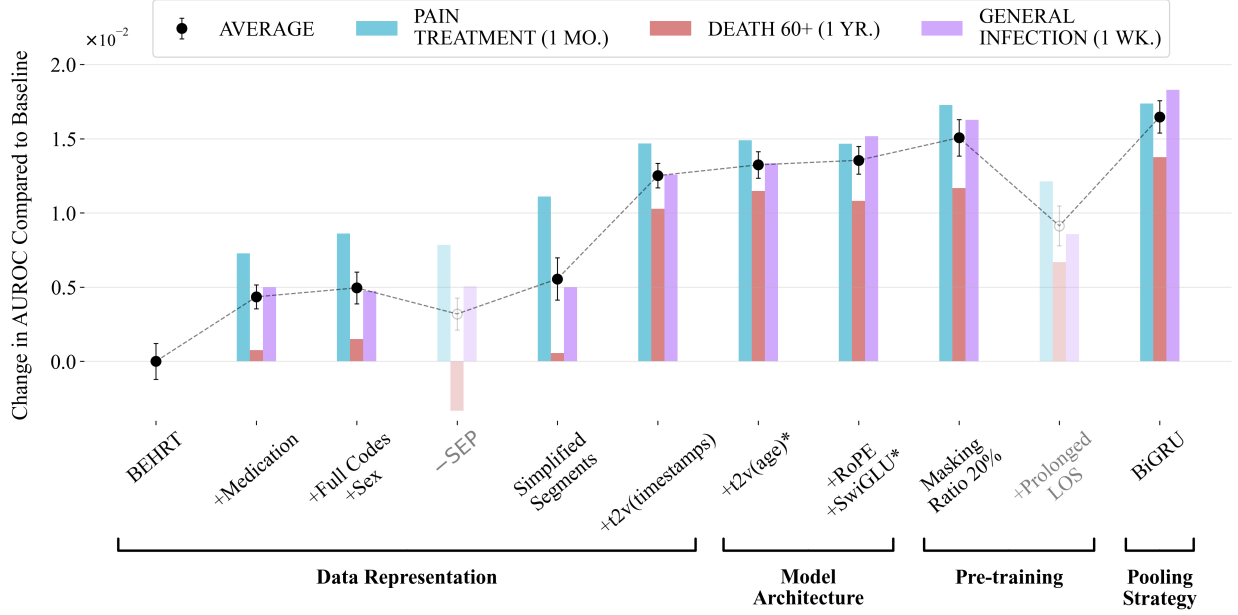


Figure 1: This figure depicts the selection process of the best configuration. Non-selected configurations are depicted in a lighter shade, indicating they were considered but not chosen based on their impact on the average Area Under the Receiver Operating Characteristic (AUROC), across three prediction tasks, whose prediction windows are denoted in parentheses. Models were trained in five-fold cross-validation and evaluated on one test set per task. An extension of these experiments for masking ratios and pooling strategies is demonstrated in Appendix I.

* Model performance metrics are averaged not only across cross-validation folds but also over multiple pre-training runs to mitigate instabilities in convergence.

metrics as seen in works like BRLTM [Y et al., 2021] and Hi-BEHRT [Li et al., 2022] or use ICU-specific data [Zhang et al., 2022]. This heterogeneity complicates direct comparisons and makes it untractable to pinpoint the sources of improvement.

The field of NLP quickly recognized that when optimally configured, base transformer models could outperform later variants [Yang et al., 2019]. A prime example is RoBERTa, which showed that an optimally configured base BERT model could achieve superior results [Liu et al., 2019]. We suspect a similar situation exists in the EHR domain, where initial approaches like BEHRT and Med-BERT may not be fully optimized.

Despite the parallels with NLP, it is also important to acknowledge the distinct characteristics of EHR data, notably its strong temporal aspect, irregular time intervals, and the distinct distribution of sequence lengths [Rasmy et al., 2021]. These disparities necessitate a comprehensive evaluation of components that have shown success in NLP within the context of EHR. Furthermore, design choices take a weighted importance in healthcare compared to NLP, as even minor performance improvements can significantly impact patient outcomes.

In contrast to fields like radiology, where deep learning models increasingly find their way into clinical practice [Yasaka and Abe, 2018, Montagnon et al., 2020, Lauritzen et al., 2023], there is a lack of trust towards EHR-based models, a cornerstone for clinical adoption [Asan et al., 2020]. Therefore, by providing insights into these models, we increase their trustworthiness and make a step towards the integration of BERT-based EHR models into clinical workflows.

Specifically, our study examines the fundamental design choices in data handling, modeling, and training processes for applying BERT-based models to EHR data, drawing inspiration from NLP optimizations. By examining model performance on a broad and diverse set of tasks, we provide robust and extensive evidence to solidify our recommendations. Furthermore, we investigate the variability in model performance to ensure that the models deployed in healthcare are optimized responsibly to inform life-critical decisions.

Main Contributions

Weak and outdated foundations, coupled with the heterogeneity in newer models, pose challenges in making direct comparisons of BERT-based EHR models. This creates a lack of trust, a major hurdle to clinical adoption, and restricts these models’ use to research settings. We address this by conducting a careful optimization rigorously evaluated across

a diverse set of tasks to increase insights into and confidence towards BERT-based EHR models. Our main contributions consist of:

- Key insights in data representation, model architecture, and training protocol through careful optimization, robustly showing model improvements
- A comprehensive and rigorous evaluation across 25 diverse clinical prediction tasks showcasing the consistency of the optimized model
- A strongly optimized and comprehensively evaluated foundation serving as a stepping stone for wider adoption of BERT-based EHR models

2 Related Work

This section examines BERT’s transformative impact on EHR analysis, illustrating its application through notable examples. We also extend our discussion to the developments related to BERT in the NLP field.

2.1 BERT in EHR

BEHRT [Li et al., 2020], serving as a proof-of-concept, was pre-trained on EHR data from 1.6 million patients, using sequences of patient events, along with age and visit details, and employed Caliber codes mapping ICD-10 (International Classification of Diseases 10th Revision) codes onto phenotypes. The model was evaluated in disease prediction tasks, showcasing superior performance over traditional recurrent neural network approaches. Their model used BERT’s [Devlin et al., 2018] self-supervised masked language modeling (MLM) pre-training (PT) technique.

Another early example was G-BERT [Shang et al., 2019], which introduced a graph-based approach that leveraged GNNs to represent the hierarchical structure of the ontologies of diagnosis and medication codes. G-BERT was trained on MIMIC-III [Johnson et al., 2016], a small but public dataset consisting of 40k patients from a critical care unit in the US. They evaluated using medication recommendation, a multilabel classification task recommending medications based on diagnosis codes.

To fully leverage the advantages of transformers in large-scale PT, Med-BERT [Rasmy et al., 2021] scaled up the PT dataset to encompass over 28 million patients. It introduced a novel PT task, aimed at predicting extended hospital stays, and expanded its vocabulary to cover all ICD-9 and ICD-10 codes. Med-BERT’s design also diverged from its predecessors by omitting age information and separator tokens and experimenting with various classification heads, among which the Bi-GRU [Schuster and Paliwal, 1997] turned out to be the most effective.

Since then, BEHRT and Med-BERT have been the foundation of many subsequent models. For example, Hi-BEHRT [Li et al., 2022] incorporated hierarchical attention mechanisms to better handle patients with extensive medical histories. CEHR-BERT [Pang et al., 2021] integrated time intervals between patient visits to address the critical aspect of temporal information in EHRs, offering a more nuanced view of patient histories. Targeted-BEHRT [Rao et al., 2022] performed causal inference on EHR data reproducing associations found in randomized controlled trials. BRLTM [Y et al., 2021] included medications, procedures, gender, and topics from clinical notes for depression and chronic disease prediction. RareBERT [Prakash et al., 2021] extended Med-BERT with an adaptive loss for class imbalance, type embeddings, and temporal reference embeddings to better handle rare diseases.

2.2 Transformers in NLP

Despite the advancements on EHR data, the journey of transformer-based models in healthcare mirrors the early challenges faced by foundational models like BERT [Devlin et al., 2018]. Initial iterations were soon recognized to be underoptimized, brought to light by replication studies such as RoBERTa [Liu et al., 2019].

Their study showed that the original BERT model was under-trained and questioned the usefulness of its Next Sentence Prediction task. They advocated for dynamic over static masking, including longer sequences, and larger batch sizes. They highlighted the importance of key design choices by showcasing better performance than later variants like XLNet [Yang et al., 2019].

We argue that this critique extends to the realm of BERT-based EHR models, which exhibit a wide range of configurations. The diversity in vocabularies, datasets, tasks, and incorporated clinical information underscores the challenge of isolating the effect of individual model components. This heterogeneity complicates the assessment of progress within the field, blurring the lines between genuine innovation and incremental modifications.

Lastly, transformer-based models have seen rapid development, primarily driven by the rise of Large Language Models (LLMs). Innovations include new activation functions, [Shazeer, 2020], improved position embeddings [Su et al., 2023], and compute efficient fine-tuning methods [Dettmers et al., 2023]. Furthermore, comprehensive advancements such as

Table 1: Statistics for the different train-val sets for death (60+), pain treatment, and general infection for the optimization experiment. We provide the prediction window and mean age at censoring as well as the prevalence of positive patients.

Condition (Age Group)	Prediction Window	Mean Age	Patients [thousands]	Prevalence [%]
Death (60+)	1 yr.	75	55.5	25.6
Pain Treatment	1 mo.	44	107.3	52.7
General Infection	7 d.	49	106.4	57.7

the transformer++ recipe have been integrated into modern LLMs like PaLM [Chowdhery et al., 2022], and Llama [Touvron et al., 2023], marking a significant shift from earlier architectures like BERT.

3 Cohort

The data consists of EHRs from 1,825 million patients, prospectively recorded during the period from 2016 to 2022 in the Capital Region of Denmark. We consider only patients with at least one code eligible, which results in 1,805 million patients. For FT, we deem at least one code to be a proper input for binary classification but use patients with ≥ 3 codes for PT to ensure a proper input for the MLM objective task. An overview of the selection pipeline can be seen in Appendix F.

3.1 Cohort Selection

For the optimization experiments, we randomly selected 500 thousand patients, split 80/20 into train and validation sets, and reserved 260 thousand for CV. The specific statistics for each CV task can be seen in Table 1. This subgroup was used as part of the PT group for the final generalizability evaluation, where we used a total of 1,202 million patients, with the remaining 702 thousand used for CV.

We observed a discrepancy in prevalence rates between the optimization and generalization experiments. This could have been caused by the assignment of the subgroup to the generalization PT group, as the subgroup was filtered, which may have left fewer positive cases for the generalization evaluation.

To test generalizability across a wide range of tasks while managing computational demands, we maintained a manageable cohort size for the main experiments. However, we also performed "**maximizing performance**" experiments, where we extended the cohorts, including PT patients in the FT set for select tasks. Details of these tasks and the expanded cohort sizes are provided in Table 3. For model selection, we require each patient in the PT cohort to have at least three diagnoses to ensure a fair comparison between the diagnosis only and the diagnosis+medication models. For the final evaluation, we require at least three of any (diagnosis or medication) codes, to accurately reflect the increased cohort when including medication.

For each condition, we defined a relevant age and/or sex and a prediction window that we use for censoring, given in Table 1 for the optimization experiments and in Appendix G for the generalization assessment. The full age distributions for the cohorts before age group selection are given in Appendix H. We consider patients to be positive when having at least one of the codes given in Appendix K and negative otherwise. We address imbalance and potential information leakage using the methods described in subsection 4.4.

3.2 Data Extraction

We extract all medical events, with an event consisting of three main variables: concept, timestamp, and admission ID. These events undergo minimal data processing through a simple three-stage process of replacing missing ICD10 codes with textual descriptions, attempting to impute missing values, only when confident, before excluding any events with missing info or ages outside a normal range. We leave a more detailed description in Appendix C.

3.3 Feature Choices

After processing, we end up with a simple data format that describes each medical event with the three main variables: concept, timestamp, and admission ID. The concept of each event is the medical code (diagnosis or medication) and is the main input to the model. We include the timestamp of the event and use it to calculate the patient’s age. The admission ID, together with the timestamp, is used to calculate the visit numbers.

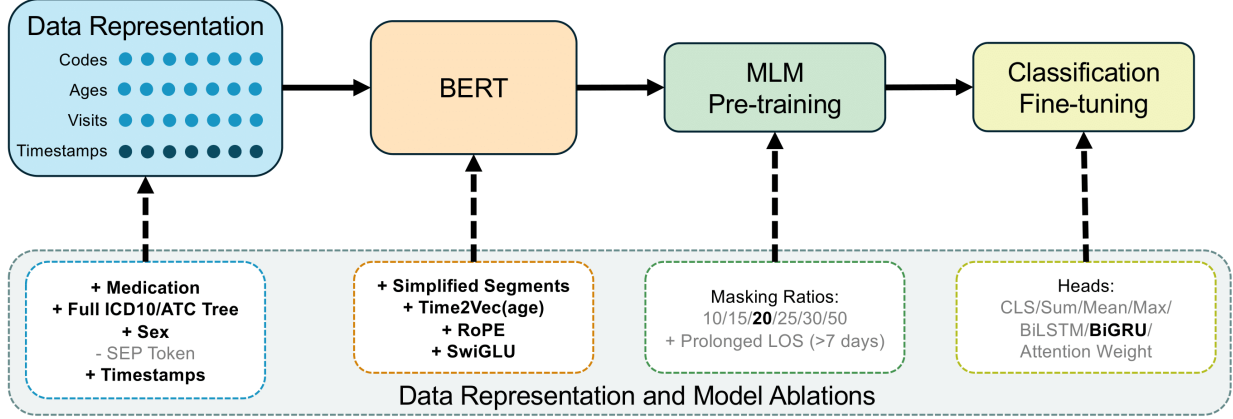


Figure 2: Overview of the iterative optimization procedure, highlighting adapted settings in bold. We augmented the input data with medication codes, patient sex, event timestamps, full International Classification of Diseases (ICD-10), and Anatomical Therapeutic Chemical (ATC) codes. Improvements to the model include integrating time2vec embeddings, Rotary Position Embeddings (RoPE), and the SwiGLU activation function. For pre-training (PT), we investigated various masking ratios and Med-BERT’s secondary PT task: predicting prolonged length of stay in the hospital. We investigated various pooling strategies for fine-tuning (FT), finalizing on Bidirectional Gated Recurrent Units (BiGRU) based on its superior results.

4 Methods

We will introduce our base model and the modifications tested, separated into a data presentation and a technical part. The data representation part investigates the impact of representing commonly available background EHR data with little to no implementation overhead. The technical part investigates the impact of improved architecture, embeddings, and training schemes, primarily relevant to machine learning practitioners. Modifications are introduced incrementally, evaluating their impact on the performance of three generic downstream tasks and including them when there is an increase in average Area Under the Receiver Operating Characteristic (AUROC). An overview of this procedure is shown in Figure 2. Next, we describe the optimization procedure before defining the various downstream tasks and close with the hyperparameters of our architecture and optimization.

4.1 Base Model

We used a slightly adapted BEHRT as our starting point. BEHRT’s input data consist of medical events with CLS (classification) and SEP (separator) tokens, age in years, and visits denoted by "positions" and binary visit encodings (segments).

In BEHRT, the events consist of level four ICD10 codes alongside Read codes [Booth, 1994], mapped to 301 Caliber codes [V et al., 2019] representing disease phenotypes. Lacking access to Read codes, we utilized level four ICD10 codes. Similarly, we employed level four ATC (Anatomical Therapeutic Chemical) codes for experiments involving medications. We believe this also strengthens the generalization of our experiments, as ICD and ATC are more commonly used than mappings such as Read and Caliber codes.

We adopted Med-BERT’s leaner configuration for hyperparameters to maintain manageable model dimensions, as BEHRT could use more layers and heads due to its smaller vocabulary size. Further details are provided in subsection 4.6. Henceforth, we will refer to our modified baseline configuration simply as BEHRT for ease of discussion.

4.2 Data Representation

We investigate different components of data representation which are easily integrated into the model and commonly ignored.

Medication Codes We adapted the BEHRT model to include medication codes, recognizing their potential to improve predictive modeling [Wei et al., 2016]. Although the inclusion of medication codes, due to their similar coding system to diagnoses, is not unconventional, it has yet to become standard practice. We consider medication a vital addition to the

model, as diagnosis codes are often incomplete or inadequate [Horsky et al., 2017]. For simplicity and computational feasibility, our current model includes only the medication codes, akin to how we handle diagnoses, without considering dosage information. Although our dataset also contains lab tests, vital signs, and procedures, with significantly higher event counts (672.2 million, 605.9 million, and 300.3 million events, respectively), we focus solely on the 227.6 million medication events due to their direct relevance to treatment decisions.

Full-Depth Codes and Background Sentence We propose to include more fine-grained data by using the full-depth medical and diagnostic codes and include the patient’s sex. It is common for EHR models to limit the input to higher-level codes, such as in BEHRT [Li et al., 2020]. This can be beneficial as high-level representation decreases the vocabulary size which in turn decreases model size. It may also reduce noise to represent information on a coarse level only. Despite these advantages, it can also remove critical information and might especially be detrimental when rare conditions are of interest. BEHRT mapped their codes to 301 Caliber categories, far less than the number of ICD10 codes. When utilizing level four codes, our model works with a vocabulary size of 2,475, whereas using the full codes expands the vocabulary to 17,469. This highlights the balance between model complexity and the granularity of information captured.

To include the patient’s sex in the model, we prepend a token representing male or female to their sequence. We foresee that other relevant demographic variables could be included as well.

Timestamps We suggest the inclusion of medical event timestamps to enhance our model, moving beyond BEHRT’s sole reliance on age for a notion of time. Timestamps allow for the precise tracking of short-term changes and are crucial for accurately predicting outcomes that depend on fine time resolution. Unlike the patient’s age, which offers a *relative* sense of time, timestamps provide an *absolute* timeline, enabling the model to account for societal shifts, including pandemics, changes in treatment protocols, and seasonal trends. This feature is particularly useful for detecting shifts in clinical EHR coding practices and adapting to variable baseline risks in conditions like influenza, which intensifies in winter.

Furthermore, external factors can affect hospital interactions, as highlighted during the COVID pandemic when many avoided hospital visits, even for severe conditions [Nourazari et al., 2021]. This shift can alter the interpretation of EHR sequences. We can account for these temporal variations in EHR semantics by incorporating event timestamps, ensuring our model remains accurate over time. Timestamps are input into the model as the time difference in hours from a specific reference point (January 26, 2020).

Separator Tokens Separator (SEP) tokens are designed to provide an additional signal to separate visits. Since visit numbers are given explicitly to the model, Med-BERT hypothesized that removing SEP tokens should not diminish performance. Furthermore, for sequences longer than the truncation length, removing SEP tokens allows for including more medical events in the sequence.

4.3 Technical Modifications

Here, we introduce the technical components, such as new embeddings and architectural changes, and pre-training (PT) objectives that were tested.

Simplified Segments We believe that BEHRT has redundant visit separation (segment) information, including separator tokens, visit numbers, and binary visit encodings. We hypothesize that the model can infer the binary visit encodings through visit numbers and differentiate between visits in this way. We remove the binary visit embedding and change BEHRT’s initialization of the visit number embedding to a standard one rather than their cosine and sinus version.

Time2Vec We suggest substituting BEHRT’s conventional age embedding with a Time2Vec [Kazemi et al., 2019] embedding. BEHRT’s current model employs embeddings that rely on discrete age in years, which can be problematic, especially in pediatric cases where changes occur on shorter timescales. These embeddings may also struggle to discern the relationship between different ages. Time2Vec models this relationship explicitly, using differently parametrized periodic functions, by giving similar representations to similar ages. Adequately understanding the relationships between ages necessitates exposure to a wide range of examples for each age. To overcome these limitations and to accommodate the timestamps mentioned previously, we propose the adoption of Time2Vec. We leave the details of Time2vec and our implementation of it in Appendix A.

Improved Transformer Recipe Since the original transformer paper, numerous refinements have been suggested. In this study, we incorporate two such advancements from the improved transformer recipe, dubbed transformer++,

that are particularly effective in boosting performance: Rotary Position Embedding (RoPE) [Su et al., 2023] and the SwiGLU activation function [Shazeer, 2020].

RoPE replaces the standard positional embedding by encoding the absolute position through a rotation matrix and incorporating relative position in self-attention. The idea is to rotate the embeddings with an angle corresponding to the positional index of the event which is applied at every key and query multiplication in each attention head.

The SwiGLU activation function is a combination of Swish [Ramachandran et al., 2017], also called Sigmoid Linear Unit (SiLU), and GLU [Dauphin et al., 2017]. It was shown to improve performance [Shazeer, 2020] and is defined as:

$$\text{SwiGLU}(x) = (\text{Swish}(xW_1) \otimes xW_2)W_3, \quad (1)$$

$$\text{Swish}(x) = x * \text{sigmoid}(x). \quad (2)$$

Pre-training Objective The training scheme of BEHRT plays a major role in the performance, as the strength of these models largely comes from PT, specifically Masked Language Modelling (MLM). Here, BEHRT copied the setup from BERT, masking 15% of the tokens—referred to as the masking ratio (MR)—using the 80-10-10 strategy (80% mask, 10% random word, 10% unchanged). We have yet to see any evidence-based reason for picking 15% in EHR, and other modalities (images, audio, and videos) have much higher MRs [Dosovitskiy et al., 2021, He et al., 2021, Baevski et al., 2020, Tong et al., 2022]. It has also been seen that larger models tend to favor larger MRs [Wettig et al., 2023].

Furthermore, we tested the second binary classification PT objective introduced by Med-BERT, in which a prediction is made on whether the patient has a prolonged length of stay in the hospital (> 7 days) for each patient. This task was run alongside MLM.

Pooling Strategies Lastly, we test different pooling strategies for the downstream tasks. It is common to use the CLS token, i.e., the first token in the sequence, as BERT introduced it as a special token specifically for classification. We investigate three additional simple pooling strategies, sum, mean, and max, which aggregate information across the sequence dimension. We also test a weighted variation of mean pooling, which uses an attention layer with a key and query to find the weights of tokens. Lastly, we test two bidirectional Recurrent Neural Networks (RNNs), BiGRU and BiLSTM, as this has been shown to increase performance [Rasmy et al., 2021].

4.4 Optimization Procedure

After introducing the various modifications, we now outline the optimization procedure in this section, as detailed in Figure 2. Our approach begins with improving the data representation, followed by model improvements, and concludes with optimizing the PT process and fine-tuning (FT) pooling strategy.

For evaluation, we use three generic downstream tasks for generalizable performance: one-year death prediction for patients older than 60 y.o., one-month pain treatment prediction, and one-week infection prediction; the corresponding definitions are given in Appendix K. Death encapsulates the final outcome and can signal the overall health status, encompassing various conditions and factors affecting patient well-being. Pain treatment is very frequent and covers a wide range of causes. It can be both mild and severe, indicative of many underlying conditions, and thus quite general. Infection points to a compromised immune system, which also tells us about the general condition of a patient.

Our three tasks constitute a generalizable evaluation by covering a broad range of health conditions. To compare the performance between models, we mainly considered the area under the receiver operating curve (AUROC) but also reported the area under the precision-recall curve (AUPRC) in Appendix J. For each downstream task, we perform five-fold cross-validation (CV) evaluating on a test set and compute the mean and standard deviation. The downstream task was a binary classification and was constructed as follows. For patients with positive outcomes, we used the outcome of interest as the index date. We then excluded all events within a specified prediction window leading up to the index date, with the window tailored to each downstream task. For patients with negative outcomes, we assign random censoring dates matching the distribution of the positive cases to prevent information leakage. To account for the imbalance in the dataset we sampled with the inverse square root of the frequency of each class [Mahajan et al., 2018, Mikolov et al., 2013].

4.5 Assessing Model Generalizability

This section outlines our methodology for evaluating the model’s performance across a diverse array of health conditions. Our goal was to ensure the generalizability of our conclusions by covering the prediction of a wide spectrum of diseases and conditions (used interchangeably from now on), showcasing the optimized model’s adaptability and precision compared to the baseline model in various clinical contexts.

Our comprehensive evaluation encompassed conditions with varying prevalence, from commonly prescribed treatments like pain medications to less frequent conditions such as pancreatic cancer. We considered how conditions develop over time ranging from chronic diseases like diabetes to acute events like major bleeding.

We also covered conditions with different associations to overall health and underlying factors. For instance, the strong correlation between stroke and overall health [Arboix, 2015] contrasts with the more complex interplay of genetic and environmental factors in schizophrenia [Stilo and Murray, 2019, Henriksen et al., 2017], which presents a subtler EHR data signature. We also covered conditions with a limited range of causes, such as anaphylaxis, which primarily involves drug-related causes, and those that can have a wide range of underlying conditions, such as sleep disorder. The conditions are listed in Appendix G together with prediction window and age groups, which were chosen based on relevance and considering the age distributions provided in Appendix H.

This structured approach to task selection is designed to demonstrate the generalizability of our model in identifying a wide range of health outcomes, underscoring its potential utility in diverse clinical scenarios without having to revisit the model choices for every new downstream task. We construct each downstream task as outlined in the previous subsection 4.4 and perform hypothesis testing as described in Appendix B.

4.6 Hyperparameters

To ensure a fair comparison, we standardized the hyperparameters across all experiments. We aimed to facilitate stable model convergence across all models while keeping manageable training durations. We utilized a base model architecture consisting of six layers and six attention heads with a hidden size of 192 and an intermediate size of 64. Sequences were truncated to 512 tokens, and a batch size of 256 for PT and 512 for FT was used. Learning rates were set to 10^{-3} for PT and $2 \cdot 10^{-5}$ for FT using a linear warmup of five epochs. Models were trained for 50 epochs during optimization and 150 epochs for the generalization experiments with an early stopping of five epochs. A more detailed overview can be found in Appendix D.

5 Results on Real Data

5.1 Evaluation Approach/Study Design

We first evaluated different components of data representation and technical modifications by pre-training for 50 epochs using a cohort of 500 thousand patients before doing a five-fold CV on death, pain treatment, and infection prediction. We iteratively added components to the pipeline if they improved the average AUROC across these three downstream tasks.

Subsequently, we compared the baseline (BEHRT) against a model with only improved data representation (BEHRT_{+D}) and a model with both improved data representation and technical components (CORE-BEHRT). We evaluated the performance on diverse tasks to ensure the generalizability of the performance gains.

Before setting up our experiments, we investigated the variance in PT and FT by running the base model multiple times. We found that both yielded stable performances with standard deviations of $6.6 \cdot 10^{-4}$ and $3.2 \cdot 10^{-4}$ for PT validation loss and FT validation AUROC, respectively. We investigated the source of variance in FT by conducting a leave-2-out CV; see Appendix E for details. The analysis revealed that the mean difference within test sets is $7.9 \cdot 10^{-4}$ and pairwise difference for the means across test sets is $4.1 \cdot 10^{-3}$, showcasing that variation is mainly attributed to the test set. These results gave us sufficient confidence to justify a single PT, followed by a five-fold CV. To ensure robust evaluation, we intentionally used larger-than-usual test sets.

5.2 Optimization Experiments

Through iterative enhancements in data representation, model architecture, and training as visualized in Figure 1, our optimization experiments have systematically advanced AUROC across multiple clinical prediction tasks.

Data Representation The proposed data changes outline two main contributions, medication, and timestamps. Both yielded a clear and significant increase in AUROC, whereas full-length codes and sex showed a smaller increase. This makes it clear that adding deeper and more precise data has a strong positive impact on the model.

However, it must be mentioned that adding medication introduced 258 new codes, which increased average sequence lengths by 143.3 codes and training time by five times compared to baseline. Changing to full-depth codes and adding sex introduced almost 15 thousand new codes, which increased the model size by 3.7 times up to 7.933 million parameters. This is due to the increased size of the embedding layer, which increases training time, as these are trainable parameters but do not affect inference time.

Notably, removing SEP tokens worsened performance, primarily for death, while the simplified segments yielded improvements. This reveals that the model mainly uses the SEP tokens to differentiate between visits and relies on explicit embeddings.

Table 2: Model performance for BEHT (baseline), BEHT_{+D} (BEHT with optimized data representation) and CORE-BEHT for multiple conditions alongside average age at censoring, patient numbers, and prevalence for positive patients in the test set. If applicable, the selected age group and sex are given in parentheses.

Condition (Age Group/Sex)	Prediction Window	Mean Age	Patients [thousands]/ Prevalence[%]	AUROC [%] BEHRT/BEHRT _{+D} / CORE-BEHRT		
Cancer						
Basal Cell Carcinoma (70+)	1 yr.	80	20.6 / 4.7	64.4	65.8**	67.2**
Breast (60+/f.)	3 mo.	74	13.5 / 4.4	57.0	57.6	62.4
Colon (60+)	3 mo.	75	35.5 / 1.3	61.4	63.6**	62.8
Lung (60+)	3 mo.	74	35.0 / 2.3	66.5	67.9*	68.4**
Pancreatic (60+)	3 mo.	74	37.7 / 0.5	52.9	57.5	56.5
Cardiovascular						
Arrhythmia (60+)	6 mo.	76	35.5 / 10.7	68.9	71.5***	71.6***
Myocardial Infarction (60+)	14 d.	75	40.4 / 2.6	65.7	68.3***	68.7***
Stroke (60+)	14 d.	75	30.2 / 1.3	64.5	69.4**	70.2***
Chronic						
Alzheimer’s Disease (70+)	3 mo.	82	23.1 / 3.7	73.9	76.5**	77.0***
Diabetes	6 mo.	65	107.4 / 4.9	83.5	84.8***	84.8***
Osteonecrosis (70+)	6 mo.	77	25.7 / 0.2	82.5	84.6	84.2
Death (60+)	1 yr.	80	35.0 / 22.7	86.6	87.5***	87.7***†
Emergencies						
Anaphylaxis	24 hr.	32	169.1 / 0.2	71.0	71.4	71.9
Diarrhea	3 d.	59	158.1 / 1.4	85.6	86.5**	87.5***†
Major Bleeding	24 hr.	59	162.6 / 4.6	78.2	78.9**	79.7***†
Infections						
General Infections	7 d.	47	141.0 / 42.8	78.2	79.7***	79.9***†
Pneumonia	14 d.	61	146.6 / 5.5	82.3	83.0***	83.2***†
Psychiatric						
Depression (15 – 30)	3 mo.	45	142.7 / 1.8	71.3	71.8**	71.9**
Schizophrenia (20 – 30)	3 mo.	24	12.7 / 1.2	83.8	85.8**	86.0**
Pain treatment	1 mo.	47	133.8 / 42.0	80.8	82.1***	82.2***††
Pregnancy-Related						
Preterm birth	at pregnancy	30	11.1 / 13.7	53.4	54.5	59.5
Fetal Growth Restriction	at pregnancy	31	11.1 / 6.9	56.9	61.6***	61.7*†††
Placental insufficiency	at pregnancy	32	11.1 / 0.5	58.8	56.3	58.6
Large fetus	at pregnancy	31	11.1 / 2.4	57.9	60.8	61.6
Others						
Sleep disorder (20 – 40)	3 mo.	30	34.0 / 4.1	74.8	75.9**	75.9**

*/**/*** Model is significantly better than BEHT with $p < 0.05/0.01/0.001$

†/††/††† Model is significantly better than BEHT_{+D} with $p < 0.05/0.01/0.001$

Table 3: Results obtained with CORE-BEHRT model for an increased cross-validation set on select conditions. AUROC is given with standard deviation computed from the five folds. The same test sets as described in Table 2 was used.

Condition (Age Group/Sex)	Patients in CV [thousands]	Prevalence [%]	AUROC [%]
Breast Cancer (60+/f.)	124.3	5.2	63.9±0.3
Preterm Birth	100.0	13.7	63.7±0.2
Schizophrenia (20 – 30)	89.2	1.4	85.9±0.6
Stroke (60+)	276.2	1.5	70.7±0.5

Technical Components The two technical components for the model, Time2Vec for age and the improved transformer recipe, yield smaller AUROC gains. We noticed during PT that RoPE, which is part of the improved transformer recipe, yielded a huge boost in MLM performance (33% \rightarrow 58% for top-1 and 82% \rightarrow 91% for top-10 accuracy). This large improvement primarily comes from an improved ability to predict masked medication codes (35% \rightarrow 61%) with a smaller improvement for diagnosis codes (20% \rightarrow 22%).

Masking Ratios We found that an MR of 20% was optimal, showing higher AUROCs for each of the three tasks compared to the base MR of 15%, as seen in Appendix I. At the same time, we observed that changes in masking ratios did not result in substantially different average AUROCs, and in general, the impact was inconsistent across the downstream tasks.

Pooling Strategies The largest improvement in pooling strategies comes from the BiGRU head, but we show in Appendix I that most techniques performed better than the standard CLS token. Even simple aggregation functions, such as mean and max pooling, which use no additional trainable parameters, showed general improvements compared to CLS.

5.3 Generalization Results

Our analysis revealed that the CORE-BEHRT model outperformed the baseline in all evaluated scenarios, as detailed in Table 2. When compared to BEHRT_{+D}, significant improvements were observed in only seven downstream tasks, notably including death, pain treatment, and general infection prediction. Despite this, the CORE-BEHRT model consistently demonstrated increased average AUROC scores across the board with exceptions in only five tasks, with two of them being equal in performance.

It is worth noting that tasks where significance was not achieved for both BEHRT_{+D} and CORE-BEHRT, compared to baseline, were predominantly among the more challenging ones, characterized by lower AUROC scores. Specifically, six out of eight tasks lacking significant improvements had AUROC scores below 0.65.

5.4 Maximizing Performance

The impact of increasing the CV set on model performance varies by condition, as seen in the comparison of Table 2 and Table 3. For breast cancer and preterm birth, performance increases are notable, rising by 2.4% and 7.1% in AUROC, respectively. In contrast, the AUROCs for stroke and schizophrenia remain largely unchanged despite the significant increases in patient numbers (3.6 times for stroke and 2.5 times for schizophrenia). This lack of improvement suggests a performance saturation at lower patient counts, supported by their already high baseline AUROCs of 70% and 86%, respectively.

6 Discussion

We have carefully optimized a BEHRT on three generic tasks, followed by a rigorous evaluation across a broad spectrum of clinical prediction tasks, showcasing the generality of our findings. This not only lays a solid foundation for future research but also sets a benchmark for the meticulous evaluation necessary to develop effective healthcare models.

The discrepancy in performance gains between the MLM and our downstream EHR tasks, observed with the improved transformer recipe, could be explained by the fact that EHR tasks are not as closely tied to the PT objective as in NLP. EHR tasks that are more closely related to MLM might minimize the discrepancy. This might also explain our observation that different masking ratios impact different tasks in an individual and inconsistent manner. Developing a

more tailored PT objective, that closely aligns with the nature of EHR downstream tasks, might alleviate this discrepancy altogether.

The generalization results show a clear trend. Improving data representation consistently, except for placental insufficiency, improves performance. However, expanding the data representation can be associated with increased model size and training times. Currently, EHR models are small compared to other domains, so increases are still manageable with a low compute budget. The technical components give the models an additional performance boost across all but four tasks and reach significance in seven of them. The technical components also never perform significantly worse and do not affect training times much, making them a solid addition to the optimized model.

Increasing the size of our FT sets to maximize performance showed limited improvements for certain tasks despite the large data increase. This suggests that performance may saturate beyond a certain number of patients, dependent on the task. This saturation might happen faster at higher base scores, as evidenced by the initial high performance for schizophrenia. This indicates that more patients might not always be the solution, but data representation and technical components might be a more promising path for these scenarios.

Here, we highlight the performance of our model on key examples to illustrate its comparative strengths. In stroke prediction, our model achieves state-of-the-art results with an AUROC of 70.7%, which exceeds the 66.9% AUROC reported in recent studies [Teoh, 2018]. It also shows comparable performance in preterm birth prediction with an AUROC of 64% versus the 65% reported in the literature [AlSaad et al., 2022], and in breast cancer prediction, where it achieves an AUROC of 63.9% nearly matching the 64.8% in existing studies [Wu et al., 2017]. Notably, we used narrower age windows in our work for stroke and breast cancer and employed fewer modalities for preterm birth. These examples illustrate our model’s strengths, even under more constrained conditions.

Our optimization experiments were conducted on three distinct tasks, introducing some potential bias. The three out of seven tasks where CORE-BEHRT significantly outperforms BEHRT_{+D} (BEHRT with optimized data representation) constituted the optimization tasks. This indicates a possibility for bias during the optimization, but we designed the tasks to be general and diverse to mitigate this effect.

Limitations Due to the sensitive nature of EHR data, it can be problematic to replicate setups exactly as public large-scale EHR datasets are rare. We thus used BEHRT with hyperparameters adapted to our dataset, which is also not publicly available. We also did not convert our codes into Caliber codes, which we believe hold greater value since ICD10 and ATC codes are more widely employed.

We found that the data representation and especially the inclusion of medication greatly impacted performance. Despite this, we only investigated the effect of medication codes and ignored sources like lab tests, vital signs, and procedure codes. These data sources are not as easy to integrate as medications and introduce many more events. This made it not only hard to integrate all data sources but also computationally infeasible for us to run a comprehensive optimization and evaluation.

A key limitation of our work is the data size and collection period. Our cohort consists of approximately 1.8 million patients, significantly smaller than the dataset used for Med-BERT and tiny compared to typical NLP data sets. The data spans only a seven-year window, limiting the number of diagnosis and medication codes, as well as the feasible prediction windows. Additionally, EHR data quality varies as delays in code registration or missing codes can occur, potentially leading to incorrect patient labeling. Although we partially address this by including free text descriptions, the issue of missing codes remains. Longer prediction windows would be preferable to avoid potential information leakage, especially for slowly developing conditions such as cancer. However, extending these windows proves challenging due to the short data collection period.

We show that improved data representation significantly boosts performance, with technical components providing an additional improvement. However, the consistent effects of these modifications remain unexplained, highlighting an important area for future research to uncover and analyze the underlying patterns.

References

- Fine MJ, Auble TE, Yealy DM, Hanusa BH, Weissfeld LA, Singer DE, Coley CM, Marrie TJ, and Kapoor WN. A prediction rule to identify low-risk patients with community-acquired pneumonia. *N Engl J Med.*, 1997.
- Mohammed Khalilia, Sounak Chakraborty, and Mihail Popescu. Predicting disease risks from highly imbalanced data using random forest. *BMC medical informatics and decision making*, 11:1–13, 2011.

- Luke Slawomirski, Luca Lindner, Katherine de Bienassis, Philip Haywood, Tiago Cravo Oliveira Hashiguchi, Melanie Steentjes, and Jillian Oderkirk. Progress on implementing and using electronic health record systems: Developments in oecd countries as of 2021. *OECD Health Working Papers*, 2023.
- Edward Choi, Mohammad Taha Bahadori, and Jimeng Sun. Doctor AI: predicting clinical events via recurrent neural networks. *CoRR*, abs/1511.05942, 2015.
- E. Choi, Andy Schuetz, Walter F. Stewart, and Jimeng Sun. Using recurrent neural network models for early detection of heart failure onset. *Journal of the American Medical Informatics Association : JAMIA*, 24:361 – 370, 2016a. URL <https://api.semanticscholar.org/CorpusID:3763233>.
- Edward Choi, Mohammad Taha Bahadori, Andy Schuetz, Walter F. Stewart, and Jimeng Sun. RETAIN: interpretable predictive model in healthcare using reverse time attention mechanism. *CoRR*, abs/1608.05745, 2016b.
- Phuoc Nguyen, Truyen Tran, Nilmini Wickramasinghe, and Svetha Venkatesh. \ mathtt {Deepr} : a convolutional net for medical records. *IEEE journal of biomedical and health informatics*, 21(1):22–30, 2016.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Yikuan Li, Shishir Rao, José Roberto Ayala Solares, Abdelaali Hassaine, Rema Ramakrishnan, Dexter Canoy, Yajie Zhu, Kazem Rahimi, and Gholamreza Salimi-Khorshidi. Behrt: transformer for electronic health records. *Scientific reports*, 10(1):1–12, 2020.
- Laila Rasmy, Yang Xiang, Ziqian Xie, Cui Tao, and Degui Zhi. Med-bert: pretrained contextualized embeddings on large-scale structured electronic health records for disease prediction. *NPJ digital medicine*, 4(1):86, 2021.
- Meng Y, Speier W, Ong MK, and Arnold CW. Bidirectional representation learning from transformers using multimodal electronic health record data to predict depression. *IEEE Journal of Biomedical and Health Informatics*, 2021.
- Yikuan Li, Mohammad Mamouei, Gholamreza Salimi-Khorshidi, Shishir Rao, Abdelaali Hassaine, Dexter Canoy, Thomas Lukasiewicz, and Kazem Rahimi. Hi-behrt: Hierarchical transformer-based model for accurate prediction of clinical events using multimodal longitudinal electronic health records. *IEEE Journal of Biomedical and Health Informatics*, 2022.
- Tianran Zhang, Muhao Chen, and Alex A.T. Bui. Adadiag: Adversarial domain adaptation of diagnostic prediction with clinical event sequences,. *Journal of Biomedical Informatics*, 2022.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. Xlnet: Generalized autoregressive pretraining for language understanding. *Advances in neural information processing systems*, 32, 2019.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach, 2019.
- Koichiro Yasaka and Osamu Abe. Deep learning and artificial intelligence in radiology: Current applications and future directions. *PLoS medicine*, 15(11):e1002707, 2018.
- Emmanuel Montagnon, Milena Cerny, Alexandre Cadrin-Chênevert, Vincent Hamilton, Thomas Derennes, André Ilinca, Franck Vandenbroucke-Menu, Simon Turcotte, Samuel Kadoury, and An Tang. Deep learning workflow in radiology: a primer. *Insights into imaging*, 11:1–15, 2020.
- Andreas D Lauritzen, My C von Euler-Chelpin, Elsebeth Lynge, Ilse Vejborg, Mads Nielsen, Nico Karssemeijer, and Martin Lillholm. Assessing breast cancer risk by combining ai for lesion detection and mammographic texture. *Radiology*, 308(2):e230227, 2023.
- Onur Asan, Alparslan Emrah Bayrak, and Avishek Choudhury. Artificial intelligence and human trust in healthcare: Focus on clinicians. *J Med Internet Res.*, 2020.
- Junyuan Shang, Tengfei Ma, Cao Xiao, and Jimeng Sun. Pre-training of graph augmented transformers for medication recommendation. *arXiv preprint arXiv:1906.00346*, 2019.
- Johnson, A., Pollard, T., Shen, and L. et al. Mimic-iii, a freely accessible critical care database. *Scientific Data*, 2016.
- Mike Schuster and Kuldip K Paliwal. Bidirectional recurrent neural networks. *IEEE transactions on Signal Processing*, 45(11):2673–2681, 1997.
- Chao Pang, Xinzhuo Jiang, Krishna S Kalluri, Matthew Spotnitz, RuiJun Chen, Adler Perotte, and Karthik Natarajan. Cehr-bert: Incorporating temporal information from structured ehr data to improve prediction tasks. In *Machine Learning for Health*, pages 239–260. PMLR, 2021.

- Shishir Rao, Mohammad Mamouei, Gholamreza Salimi-Khorshidi, Yikuan Li, Rema Ramakrishnan, Abdelaali Hassaine, Dexter Canoy, and Kazem Rahimi. Targeted-behrt: Deep learning for observational causal inference on longitudinal electronic health records. *IEEE Transactions on Neural Networks and Learning Systems*, 2022.
- P. Prakash, S. Chilukuri, N. Ranade, and S. Viswanathan. Rarebert: Transformer architecture for rare disease patient identification using administrative claims. *Proceedings of the AAAI Conference on Artificial Intelligence*, 2021.
- Noam Shazeer. Glu variants improve transformer, 2020.
- Jianlin Su, Yu Lu, Shengfeng Pan, Ahmed Murtadha, Bo Wen, and Yunfeng Liu. Roformer: Enhanced transformer with rotary position embedding, 2023.
- Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. Qlora: Efficient finetuning of quantized llms, 2023.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Ben Hutchinson, Reiner Pope, James Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Levskaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier Garcia, Vedant Misra, Kevin Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayanan Pillai, Marie Pellat, Aitor Lewkowycz, Erica Moreira, Rewon Child, Oleksandr Polozov, Katherine Lee, Zongwei Zhou, Xuezhi Wang, Brennan Saeta, Mark Diaz, Orhan Firat, Michele Catasta, Jason Wei, Kathy Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov, and Noah Fiedel. Palm: Scaling language modeling with pathways, 2022.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. Llama: Open and efficient foundation language models, 2023.
- Nick Booth. What are the read codes? *Health libraries review*, 11(3):177–182, 1994.
- Kuan V, Denaxas S, Gonzalez-Izquierdo A, Direk K, Bhatti O, Husain S, Sutaria S, Hingorani M, Nitsch D, Parisinos CA, Lumbers RT, Mathur R, Sofat R, Casas JP, Wong ICK, Hemingway H, and Hingorani AD. A chronological map of 308 physical and mental health conditions from 4 million individuals in the english national health service. *Lancet Digit Health*, 2019.
- Wei-Qi Wei, Pedro L Teixeira, Huan Mo, Robert M Cronin, Jeremy L Warner, and Joshua C Denny. Combining billing codes, clinical notes, and medications from electronic health records provides superior phenotyping performance. *Journal of the American Medical Informatics Association*, 23(e1):e20–e27, 2016.
- Jan Horsky, Elizabeth A Drucker, and Harley Z Ramelson. Accuracy and completeness of clinical coding using icd-10 for ambulatory visits. In *AMIA annual symposium proceedings*, volume 2017, page 912. American Medical Informatics Association, 2017.
- Sara Nourazari, Samuel R Davis, Rachel Granovsky, Randolph Austin, Dean J Straff, Joshua W Joseph, and Leon D Sanchez. Decreased hospital admissions through emergency departments during the covid-19 pandemic. *The American journal of emergency medicine*, 42:203–210, 2021.
- Seyed Mehran Kazemi, Rishab Goel, Sepehr Eghbali, Janahan Ramanan, Jaspreet Sahota, Sanjay Thakur, Stella Wu, Cathal Smyth, Pascal Poupart, and Marcus Brubaker. Time2vec: Learning a vector representation of time, 2019.
- Prajit Ramachandran, Barret Zoph, and Quoc V. Le. Searching for activation functions, 2017.
- Yann N. Dauphin, Angela Fan, Michael Auli, and David Grangier. Language modeling with gated convolutional networks, 2017.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale, 2021.
- Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners, 2021.
- Alexei Baevski, Henry Zhou, Abdelrahman Mohamed, and Michael Auli. wav2vec 2.0: A framework for self-supervised learning of speech representations, 2020.
- Zhan Tong, Yibing Song, Jue Wang, and Limin Wang. Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training, 2022.
- Alexander Wettig, Tianyu Gao, Zexuan Zhong, and Danqi Chen. Should you mask 15

- Dhruv Mahajan, Ross B. Girshick, Vignesh Ramanathan, Kaiming He, Manohar Paluri, Yixuan Li, Ashwin Bharambe, and Laurens van der Maaten. Exploring the limits of weakly supervised pretraining. *CoRR*, abs/1805.00932, 2018. URL <http://arxiv.org/abs/1805.00932>.
- Tomás Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. Distributed representations of words and phrases and their compositionality. *CoRR*, abs/1310.4546, 2013. URL <http://arxiv.org/abs/1310.4546>.
- Adrià Arboix. Cardiovascular risk factors for acute stroke: Risk profiles in the different subtypes of ischemic stroke. *World Journal of Clinical Cases: WJCC*, 3(5):418, 2015.
- Simona A Stilo and Robin M Murray. Non-genetic factors in schizophrenia. *Current psychiatry reports*, 21:1–10, 2019.
- Mads G Henriksen, Julie Nordgaard, and Lennart B Jansson. Genetics of schizophrenia: overview of methods, findings and limitations. *Frontiers in human neuroscience*, 11:322, 2017.
- Douglas Teoh. Towards stroke prediction using electronic health records. *BMC medical informatics and decision making*, 18:1–11, 2018.
- Rawan AlSaad, Qutaibah Malluhi, and Sabri Boughorbel. Predictptb: an interpretable preterm birth prediction model using attention-based recurrent neural networks. *BioData Mining*, 15(1):6, 2022.
- Yirong Wu, Elizabeth S Burnside, Jennifer Cox, Jun Fan, Ming Yuan, Jie Yin, Peggy Peissig, Alexander Cobian, David Page, and Mark Craven. Breast cancer risk prediction using electronic health records. In *2017 IEEE International Conference on Healthcare Informatics (ICHI)*, pages 224–228. IEEE, 2017.

A Time2vec Implementation

Given a scalar input τ Time2Vec (t2v) produces a vector of size k as follows:

$$\text{t2v}(\tau)[i] = \begin{cases} \omega_i \tau + \phi_i, & \text{if } i = 0. \\ \mathcal{F}(\omega_i \tau + \phi_i), & \text{if } 1 \leq i \leq k, \end{cases} \quad (3)$$

where $\text{t2v}(\tau)[i]$ is the i^{th} element of $\text{t2v}(\tau)$, \mathcal{F} is the periodic function, where we choose cosine, ω_i and ϕ_i are learnable frequency and phase-shift parameters. Time2vec uses fewer parameters than a standard embedding and can easily be implemented as a drop-in replacement.

We noticed a certain degree of instability when using time2vec out of the box, but observed that scaling τ and clipping the 0^{th} dimension helped with this issue and stabilized the performance. After a rough hyperparameter search, we found stability when scaling age with a factor of 10^{-2} , constraining it from 0 to around 1, and scaling absolute time with a factor of 10^{-4} , converting a year of absolute time (8760 hours) to ≈ 1 . Furthermore, the 0^{th} dimension was clipped to $[-100, 100]$ to avoid large embedding values for very early events.

B P-value Computation

To assess the statistical significance between the base and best models' performances, we calculated one-sided p-values using a t-test, under the null hypothesis that the model is not superior to either BEHT or BEHT_{+D}, if applicable. The means and standard deviations, based on five observations each, allowed for precise standard error estimation and subsequent t-statistic computation for each disease category. We apply the Benjamini-Hochberg procedure to account for the false discovery rate.

C Data processing

Code-level replacement Some events had missing ICD10 codes, but also an accurate textual description of the diagnosis. Rather than excluding the events with missing ICD10 codes, we replace the missing value with the textual description. We account for these replacements when defining conditions, as each diagnosis code can now have multiple tokens to represent it.

Admission-level Imputation Missing admission IDs and timestamps can sometimes be inferred by using surrounding events. Specifically, we try to impute admission IDs and timestamps, *only* in the places where this can be done confidently.

Missing admission IDs can be inferred when a missing admission ID exists between two identical admissions, e.g. [A None A]. Here, it is highly likely that the missing admission is part of the A admissions. Any missing admission ID that cannot be inferred is replaced with a unique identifier and treated as a standalone admission.

Some missing timestamps can be imputed if they are part of an admission with at least one other entry. Here, the missing timestamp can be estimated as most admissions lie within a similar time frame. The first, average, or last timestamp can be used, where we choose the last so as not to interfere with the ordering of the other codes.

Event-level Exclusion Any event still containing missing values or events where the patient's age is outside a normal range ($0 \leq \text{age} \leq 120$) is removed.

D Hyperparameters

Architecture Specification As part of our architecture choices, we adopted settings from Med-BERT with six layers, six attention heads, and a hidden dimension size of 192. These were chosen due to their efficient performance on larger datasets. The intermediate size is 64, with a context length of 512, considering variations in vocabulary sizes which affect model size.

Optimization and Training Parameters We used an epsilon value of 10^{-6} for the AdamW optimizer. The learning rates were set to 10^{-3} for PT and $2 \cdot 10^{-5}$ for FT, with a linear warmup phase of five epochs. Optimization experiments were limited to 50 epochs with an early stopping mechanism based on validation loss for PT and AUROC for FT. For the generalization experiments, the number of training epochs was increased to 150.

Computational Resources and Training Duration Training was conducted using Nvidia A100 GPUs with 80 GiB of VRAM. Batch sizes were set at 256 for PT and 512 for FT, ensuring maximal GPU utilization for our largest model versions. The duration for PT one model on 400,000 patients ranged from 2.5 to 16.5 hours, depending on model size.

E Leave-2-Out Cross Validation

The main objective of the leave-2-out CV experiment was to analyze the sources of variance in model performance. The dataset was divided into five distinct folds, facilitating a comprehensive assessment through multiple test runs. During each iteration of the experiment, two of the five folds were selected: one served as the validation set and the other as the test set. This configuration allowed each test set to be evaluated by two different models, which were trained on overlapping subsets of the remaining three folds. The experiment was structured to execute a total of 10 unique runs, each providing insights into model behavior under different training and validation scenarios. For each test set, we computed the performance differences between the two models and then averaged these differences to quantify the variability within the model evaluations. Furthermore, this average was compared against the mean of the pairwise differences in the mean performances across different test sets. This comparison aimed to distinguish between the intrinsic variability of the model and variability attributable to the composition of the test sets.

F Cohort Selection

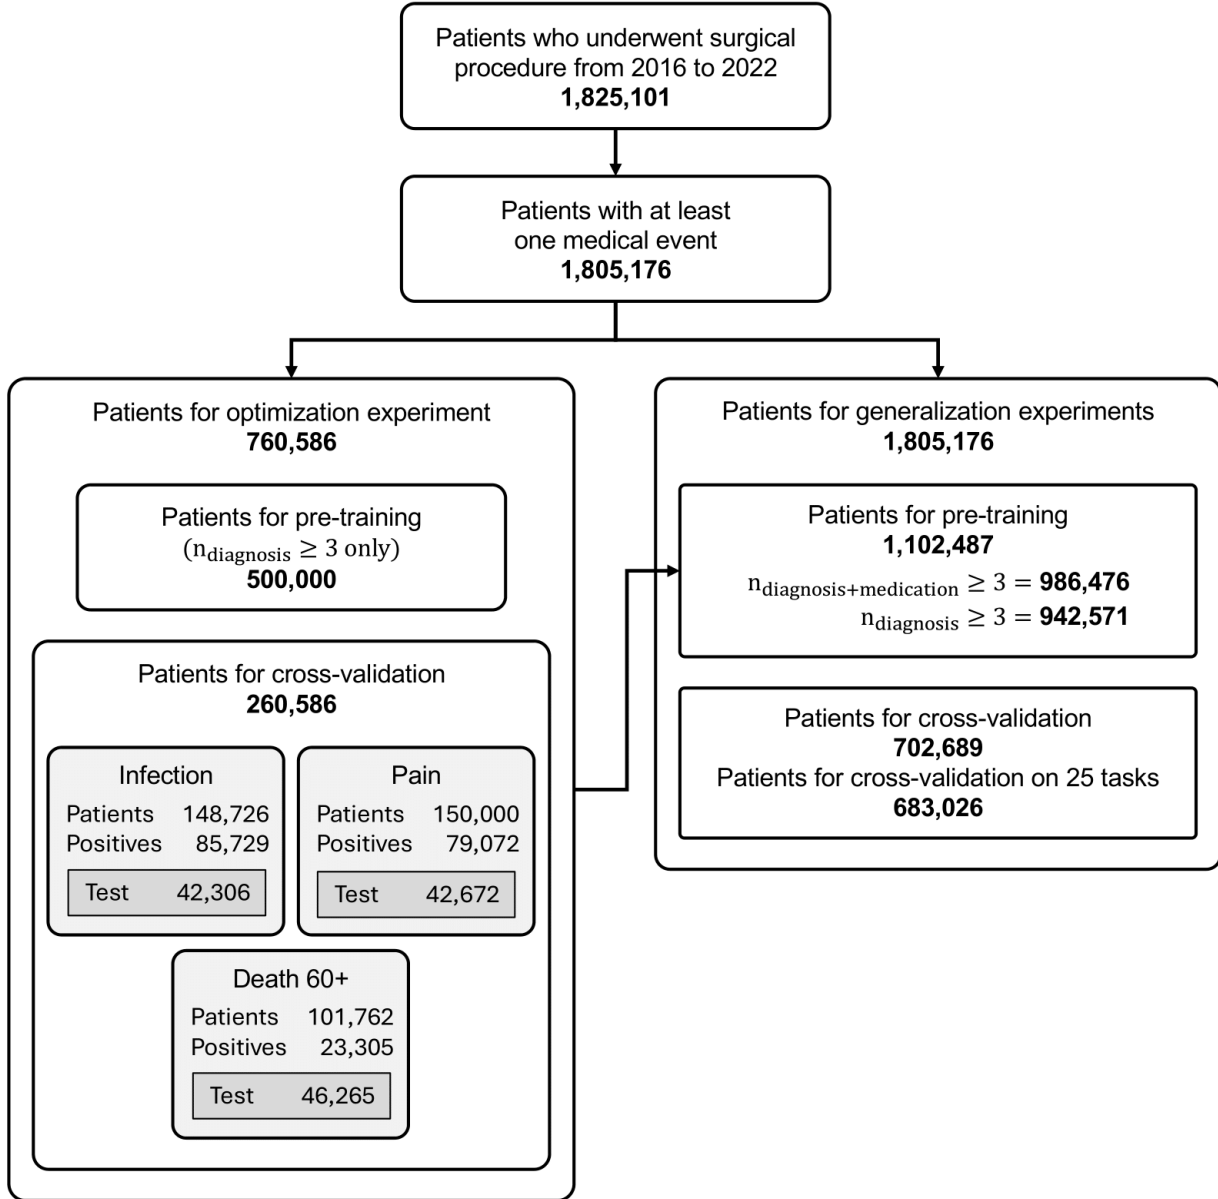


Figure 3: Selection pipeline for model selection and final evaluation

G Generalization Cohort Statistics

Table 4: Statistics for all 25 CV sets for the generalization experiment. sequence length is N events per patient sequence (post-censoring). Trajectory length* is in years;

Condition	Prediction Window	Age mean \pm std	Sequence Length Q2(Q1-Q3)	Trajectory Length Q2(Q1-Q3)	Patients [thousands]/ Prevalence[%]
Alzheimer's Disease	3 mo.	79.1 \pm 7.1	27 (3-161)	2.5 (1.3-3.9)	46.2 / 3.7
Anaphylaxis	24 hr.	41.6 \pm 24.7	4 (1-28)	2.2 (1.0-3.5)	338.1 / 0.2
Arrhythmia	6 mo.	73.6 \pm 9.2	9 (2-85)	2.4 (1.4-3.7)	70.9 / 10.7
Basal Cell Carcinoma	1 yr.	79.4 \pm 7.2	22 (3-147)	3.0 (1.9-4.3)	41.2 / 4.6
Breast Cancer	3 mo.	74.4 \pm 9.6	11 (2-102)	2.3 (1.2-3.8)	27.1 / 4.5
Colon Cancer	3 mo.	73.7 \pm 9.3	14 (2-110)	2.4 (1.3-3.9)	71.0 / 1.3
Lung Cancer	3 mo.	73.7 \pm 9.2	14 (2-110)	2.4 (1.3-3.9)	69.9 / 2.1
Pancreatic Cancer	3 mo.	73.7 \pm 9.3	15 (3-117)	2.5 (1.3-3.8)	75.5 / 0.4
Death	1 yr.	73.7 \pm 9.1	11 (2-89)	3.0 (2.0-4.3)	70.1 / 22.5
Depression	3 mo.	22.5 \pm 4.3	2 (1-7)	2.1 (1.1-3.4)	53.5 / 2.7
Diabetes	6 mo.	41.9 \pm 24.7	3 (1-17)	2.0 (1.1-3.1)	214.8 / 5.0
Diarrhea	3 d.	41.5 \pm 24.7	4 (1-26)	1.9 (0.9-3.1)	316.3 / 1.3
Infection	7 d.	41.2 \pm 24.7	2 (1-7)	1.4 (0.6-2.7)	281.9 / 42.6
Major Bleeding	24 hr.	41.7 \pm 24.7	4 (1-26)	2.0 (0.9-3.4)	325.1 / 4.7
Myocardial Infarction	14 d.	73.7 \pm 9.3	16 (3-118)	2.3 (1.1-3.7)	80.8 / 2.7
Osteonecrosis	6 mo.	79.4 \pm 7.2	32 (4-176)	3.0 (1.8-4.4)	51.4 / 0.2
Pain	1 mo.	40.8 \pm 24.5	2 (1-5)	1.4 (0.6-2.5)	267.5 / 42.0
Treatment					
Pneumonia	14 d.	41.7 \pm 24.8	4 (1-24)	1.8 (0.8-3.3)	293.2 / 5.5
Large Fetus	at pregnancy	30.4 \pm 5.1	3 (2-7)	0.5 (0.0-2.1)	22.2 / 2.3
Placental Insufficiency	at pregnancy	30.4 \pm 5.0	3 (2-7)	0.4 (0.0-2.1)	22.2 / 80.6
Fetal Growth Disorder	at pregnancy	30.4 \pm 5.1	3 (2-7)	0.4 (0.0-2.1)	22.2 / 6.6
Preterm Birth	at pregnancy	30.4 \pm 5.0	3 (2-7)	0.4 (0.0-2.1)	22.2 / 13.7
Schizophrenia	3 mo.	24.7 \pm 2.9	2 (1-7)	1.6 (0.8-2.9)	25.5 / 1.2
Sleep Disorder	3 mo.	29.3 \pm 5.6	2 (1-11)	1.8 (1.0-2.9)	67.9 / 4.3
Stroke	14 d.	73.5 \pm 9.1	12 (2-102)	1.9 (0.9-3.5)	60.4 / 1.4

* From first encounter to censoring

H Age Distribution for Generalization Cohort

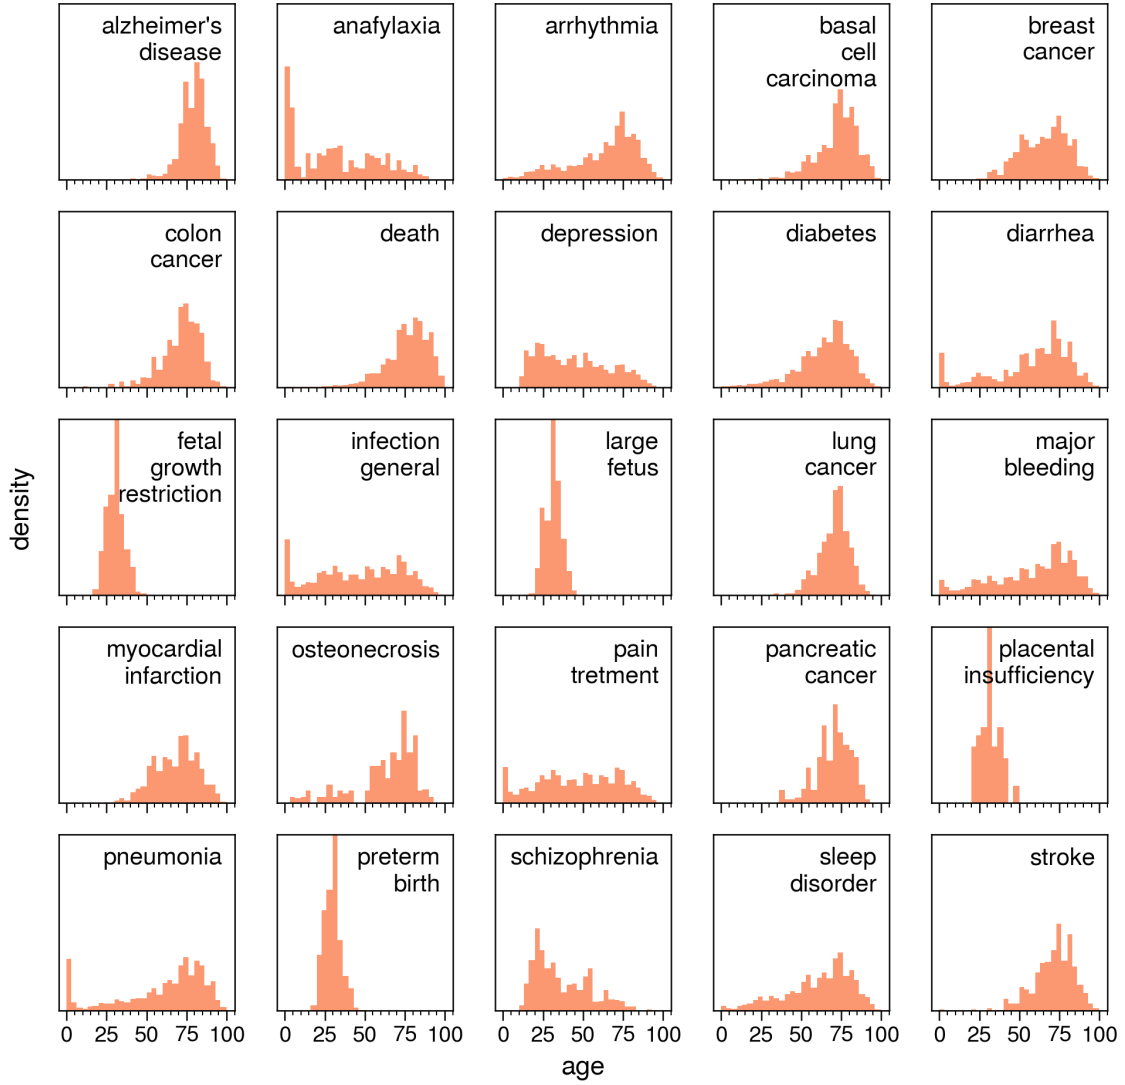


Figure 4: Age distribution for positive patients in the cross-validation set, before age group selection, for each of the 25 downstream tasks.

I Results of Generalization Experiments

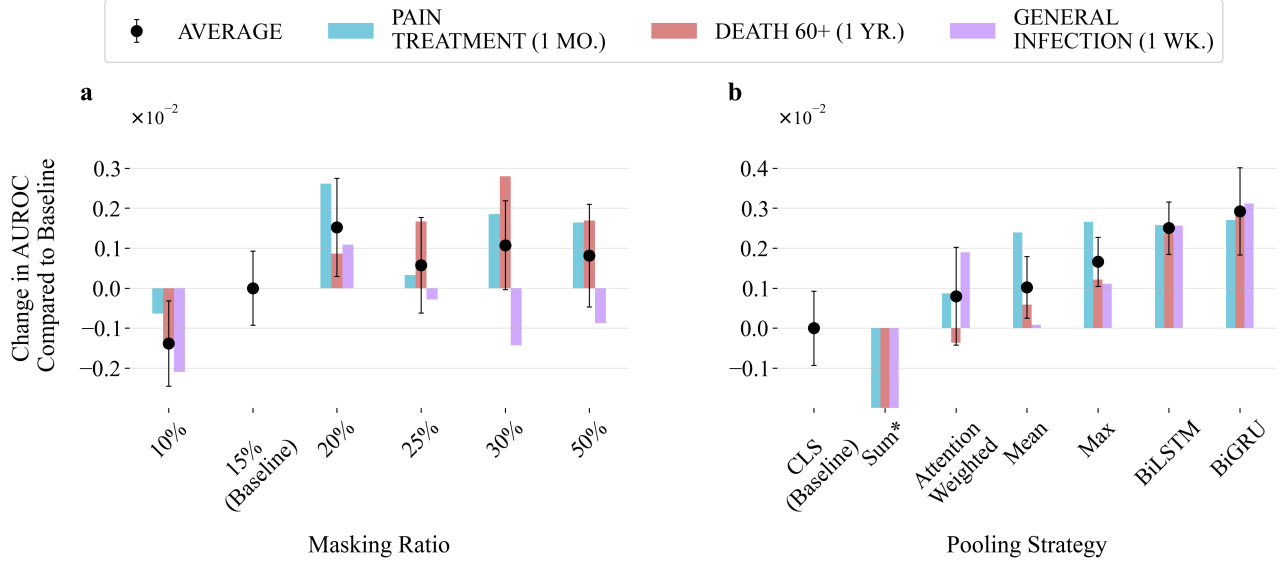


Figure 5: Ablation studies analyzing the effects of different masking ratios (MR) and pooling strategies on predictive performance. In (a) we compare the change in AUROC across varying MRs using CORE-BEHRT with 15% MR and CLS (classification) token pooling as a baseline. While an MR of 20% shows the highest average AUROC improvement, certain conditions such as death prediction benefit from a higher MR of 30%. In (b) we evaluate various pooling strategies against CORE-BEHRT with CLS pooling as a baseline. The Bidirectional Gated Recurrent Unit (BiGRU) pooling strategy emerges as superior based on the average AUROC improvement across conditions.

* This change resulted in a change of average AUROC of $-1.5 \cdot 10^{-2}$, which is not displayed to maintain readability of other differences.

J Results of Optimization Experiments

Table 5: Test scores for initial experiments given as AUROC and AUPRC with mean and standard deviations obtained from the five models trained in CV for three prediction tasks pain treatment, death, and general infection. Our intermediate models are highlighted in bold. Adapted settings are shown with an asterisk.

Model	AUROC (%)			AUPRC (%)		
	Pain	Death	Infection	Pain	Death	Infection
Data representation						
Med-BERT	76.53±0.07	79.23±0.13	72.67±0.11	79.10±0.05	52.77±0.21	77.92±0.13
*BEHRT	77.01±0.13	84.54±0.15	73.87±0.07	79.51±0.17	60.27±0.35	78.84±0.06
*+Medication	77.74±0.08	84.62±0.06	74.37±0.10	80.28±0.07	60.91±0.23	79.50±0.09
*+FV+Sex	77.87±0.06	84.69±0.17	74.34±0.04	80.44±0.06	61.24±0.37	79.39±0.06
-SEP	77.80±0.02	84.21±0.17	74.37±0.07	80.30±0.05	61.13±0.39	79.34±0.06
*Simplified Segments	78.12±0.04	84.60±0.16	74.37±0.19	80.50±0.04	61.06±0.29	79.30±0.19
*+t2v(timestamps)	78.48±0.11	85.57±0.11	75.12±0.07	80.99±0.06	63.41±0.30	80.29±0.10
BEHRT_{+D}						
Technical components						
*+t2v(age)	78.50±0.05	85.69±0.11	75.20±0.06	80.97±0.04	63.67±0.27	80.33±0.10
*+RoPE+SwiGLU	78.48±0.05	85.62±0.13	75.38±0.08	80.81±0.07	63.56±0.42	80.60±0.13
Masking Ratio 10%	78.41±0.06	85.48±0.13	75.21±0.12	80.82±0.07	63.44±0.23	80.27±0.17
*Masking Ratio 20%	78.74±0.12	85.77±0.08	75.57±0.13	81.08±0.09	63.89±0.23	80.62±0.10
Masking Ratio 25%	78.64±0.06	85.94±0.15	75.24±0.11	80.96±0.08	64.22±0.19	80.44±0.09
Masking Ratio 30%	78.56±0.06	85.74±0.19	75.41±0.11	80.89±0.03	63.94±0.42	80.42±0.14
Masking Ratio 50%	78.23±0.02	85.21±0.20	74.72±0.12	80.56±0.02	63.28±0.32	79.78±0.12
MR 20% + PLOS	78.66±0.13	85.68±0.08	75.69±0.10	80.94±0.15	63.44±0.45	80.65±0.13
Max Pooling	78.71±0.04	85.74±0.06	75.50±0.09	81.12±0.03	64.13±0.15	80.66±0.09
Mean Pooling	78.66±0.06	85.94±0.14	75.69±0.11	81.12±0.03	63.94±0.33	80.89±0.11
Sum Pooling	78.16±0.15	83.45±1.10	73.37±1.04	80.24±0.21	56.87±2.03	78.05±0.90
Attention Weighted	78.84±0.08	85.86±0.08	75.74±0.10	81.22±0.09	63.63±0.13	80.84±0.11
BiLSTM	78.74±0.06	85.74±0.08	75.54±0.06	81.18±0.09	64.22±0.20	80.62±0.09
*BiGRU	78.69±0.13	85.94±0.14	75.70±0.09	81.22±0.11	63.94±0.33	80.89±0.09
CORE-BEHRT						

K Definitions of Conditions

Table 6: Table of conditions and their corresponding codes as given in the SKS-browser, sorted alphabetically by condition.

Condition	Codes
Alzheimer's Disease	DG30, DF00
Anaphylaxis	DT782, DT788, DT789, DT882, DT886, DT881
Arrhythmia	DI47, DI49, DR00, MC01A, MC01B, Abnorm hjerterytme
Basal Cell Carcinoma	DC44, anden hudkræft exclude: Unspecified Cancer
Breast Cancer	DC50, brystkræft exclude: Unspecified Cancer
Colon Cancer	DC18, kræft i tyktarmen exclude: Unspecified Cancer
Depression	Depression, DF32, DF33
Diabetes	MA10BA02, MA10BD07, MA10BD08, MA10BD10, MA10BD11, MA10BD13, MA10BD19, MA10BD21, MA10BD24, MA10BH01, MA10BH02, MA10BH03, MA10BH04, MA10BH05, MA10A, MA10BB, MA10BK, MA10BJ, DE10, DE11, diabetes
Diarrhea	DK529, DK521, MA07D
Fetal Growth Restriction	DO365
Fetal Macrosomia	DO366
Lung Cancer	DC34, Kræft i bronkier eller lunge med metastaser, Kræft i lungens, kræft i lunge exclude: Unspecified Cancer
Major bleeding	DD62, DI60, DI61, DI62, DI690, DI691, DJ942, DK250, DK252, DK254, DK256, DK260, DK262, DK264, DK276, DK270, DK272, DK274, DK276, DK280, DK282, DK284, DK286, DK290, DK298A, DK625, DK638C, DK920, DK921, DK922, DN02, DN93, DR04, DR31, DS064, DS065, DS066
Myocardial Infarction	DI21, Komplikationer i efterforløbet af akut myokardieinfarkt
Osteonecrosis	DM87
Pain Treatment	MN02A, MN02B, MM02A
Pancreatic Cancer	DC25, Kræft i caput pancreatis med metastaser, Kræft i cauda pancreatis, Kræft i corpus pancreatis exclude: Unspecified Cancer
Placental Insufficiency	DO365A
Pregnancy	DZ340, DZ348, DZ349
Preterm Birth	DO42, DO47, DO60
Schizophrenia	DF20
Sleep Disorder	DF51, DG47, MN05CF, MN05CH
Stroke	DI64, Slagtilfælde uden oplysning om blødning eller infarkt
Unspecified Cancer	Kræftsygdom UNS med metastaser, Kræft UNS overgribende flere lokalisationer med metastaser

Table 4 (continued)

Condition	Codes
Infection	<p>DA, DB, DD738C, DD738D, DD762, DE060B, DG00-DG07, DG531, DG630, DG734, DG940, DI00-DI02, DI30, DI320, DI321, DI33, DI38, DI39, DI400-DI412, DI430, DI520, DI521, DK112, DK113, DK12, DK140, DK230, DK770, DK858E, DK930, DM00, DM01, DM03, DM461-DM465, DM490-DM493, DM600, DM630, DM632, DM650, DM651, DM680, DM710, DM711, DM726, DM730, DM731, DM86, DM900-DM902, DN61, DN70-DN77, MJ01 DN080, DN109A-DN109C, DN12, DN136, DN151, DN160, DN288D, DN288E, DN288F, DN290, DN291, DN30, DN330, DN340, DN341, DN342A, DN342B, DN390, DN412, DDN45, DN481, DN482, DN49, DN511, DN512 DA0, DK35, DDK37, DK57, DK65, DK67, MJ01, DH60, DH610, DH620, DH621, DH623, DH65, DH66, DH670, DH671, DH68, DH70, DH730, DH750, DH830, DH940 DA46, DL00-DL08, DL303, DL738H, DL88 DA541, DB43, DD733, DE060A, DE236A, DE321, DG06, DG07, DH000A, DH050A, DH440A, DH600, DJ340A, DJ36, DJ383D, DJ387G, DJ390, DJ391, DJ398A, DJ851, DJ852, DJ853, DK113, DK122, DK130A, DK140A, DK209A, DK353A, DK353B, DK570, DK572, DK574, DK578, DK61, DK630, DK650, DK750, DK810A, DK858A, DL02, DL050, DL059, DM608A, DM868A, DM869A, DN151, DN340, DN412, DN450, DN482, DN492A, DN619A, DN619B, DN700A, DN700B, DN710A, DN730A, DN730B, DN732A, DN732B, DN733A, DN735A, DN738A, DN738C, DN751, DN764, DN768A, 'infektjose', 'infektion',</p> <p>exclude: DA06, DA07, DA33, DA630, DA65, DA66, DA67, DA68, DA69, DA881, DB079D, DB18, DB35, DB36, DB37, DB38, DB39, DB4-DB7, DB80-DB83, DB85-DB92, DB94, DI300, DI521C, DK112B, DM031A, DM036, DM863, DM864, DM865, DM866 DN701, DN702, DN711, DN729B, DN729E, DN731, DN734, DN736, DN738B, DN75, DN750, DN758, DN758B, DN759, DN761A, DN761C, DN761E, DN763A, DN765, DN766 DN301, DN302, DN304, DN308, DN459B, DN481D DA06, DA07, DK573C, DK573D, DK574, DK574A, DK575, DK578, DK579, DK658A, DK658C, DK658E, DK658F, DK658G, DK658H DH604, DH604A, DH605, DH605B, DH608, DH608A, DH652, DH653, DH654, DH654C, DH661, DH662, DH663, DH681, DH701, DH708 DJ340E, DJ340F, DJ340G, DJ340H, DL889, DL889C, DA541B, DB430, DB438, DB439, DK570B, DK570C, DK572B, DK572C, DK574A, DK650M, DK650N, DK650O, DK650P</p>