

Continuous Predictive Modeling of Clinical Notes and ICD Codes in Patient Health Records

Mireia Hernandez Caralt Clarence Boon Liang Ng Marek Rei

Imperial College London, United Kingdom

{mireia.hernandez-caralt22,clarence.ng21,marek.rei}@imperial.ac.uk

Abstract

Electronic Health Records (EHR) serve as a valuable source of patient information, offering insights into medical histories, treatments, and outcomes. Previous research has developed systems for detecting applicable ICD codes that should be assigned while writing a given EHR document, mainly focusing on discharge summaries written at the end of a hospital stay. In this work, we investigate the potential of predicting these codes for the whole patient stay at different time points during their stay, even before they are officially assigned by clinicians. The development of methods to predict diagnoses and treatments earlier in advance could open opportunities for predictive medicine, such as identifying disease risks sooner, suggesting treatments, and optimizing resource allocation. Our experiments show that predictions regarding final ICD codes can be made already two days after admission and we propose a custom model that improves performance on this early prediction task.

1 Introduction

Electronic health records (EHR) are rich repositories of patient information, chronicling their medical history, diagnoses, treatment plans, medications and outcomes (Jensen et al., 2012; Johnson et al., 2016). The aggregation and modeling of this data over time presents a unique opportunity for revealing patterns that can improve patient care, operational efficiency, and healthcare delivery. Contained within the EHR are textual notes written by clinicians during patient encounters, which are essential for a comprehensive understanding of patient health. These free-text narratives stand out as a particularly rich source of nuanced information, but their unstructured format and domain-specific language use have left them largely underutilized, compared to more readily available structured data sources (Tayefi et al., 2021).

Category	# notes	# words
Discharge summ.	59,652	79,649,691
ECG	209,051	5,625,393
Echo	45,794	12,810,062
Nursing	1,046,053	185,856,841
Physician	141,624	322,961,183
Radiology	522,279	165,805,982
Respiratory	31,739	11,957,187
Other	26,988	13,086,023

Table 1: The number of clinical notes from different categories, along with the number of words in those notes, in the MIMIC-III dataset.

Health records are often accompanied by the International Classification of Diseases (ICD) codes – standardized codes that categorize diagnoses and procedures performed during clinical encounters (Cartwright, 2013). Assigning ICD codes manually is a highly time-consuming task necessary for billing, therefore previous research has been developing multi-label classification systems to detect applicable codes that should be assigned while writing a given document (Mullenbach et al., 2018; Liu et al., 2021). The research focus has been on classifying discharge summaries, which are written at the end of a patient’s hospital stay (Ji et al., 2021; Dai et al., 2022). While this setup provides a useful proxy task, the complete EHR sequence is much longer, containing detailed reports from nursing and radiology, along with specialized notes on echographies and cardiograms (FCG), among others (Table 1). Recent work has argued that for most practical applications such code classification should be performed on earlier medical notes instead of the discharge summary (Cheng et al., 2023).

In this work, we investigate the potential of predicting ICD codes for the whole patient stay at different time points during their stay. Beyond the

task of detecting codes for a given note, we treat ICD codes as a structured summary of all the treatments provided and diagnoses assigned during a hospital stay. The development of models for predicting this information early in the clinical timeline, based on partial indicators even before these codes have been officially assigned, would open many possibilities for predictive medicine. Such systems would go beyond the post-discharge diagnostic practices and could be used for identifying early disease risks, suggesting potential treatments or optimizing hospital resource allocation.

We investigate the feasibility of this novel task and evaluate to what extent the final set of ICD codes can be predicted at earlier stages during the hospital stay. In addition, we propose a custom model for this task that is able to improve prediction accuracy at different time steps. Unlike previous ICD code prediction models, the architecture is designed with causal attention to ensure that representations at any point throughout a patient’s hospital stay are constructed based on the notes available up to that point, without accessing information in the future. The model is then optimized to predict ICD codes after every additional note in the input sequence, instead of only at the discharge summary, teaching it to make predictions at any chosen time point during the hospital stay. This task poses additional challenges, as the length of the complete EHR sequence far exceeds that of the discharge summary and early notes have a weaker correlation with the final labels. We introduce a novel method that both augments the data during training and extends the context during inference, substantially enhancing the performance on early ICD code prediction. The code for the model and the experiments are available online.¹

2 Related Work

The closest previous research to ours has been on automating ICD code assignment. Given a document mentioning diagnoses and treatments in free-form text, the aim is to detect the correct codes that should be assigned by the clinician. The first attempts at this task primarily relied on convolutional neural networks (CNNs) (Mullenbach et al., 2018; Li and Yu, 2020; Liu et al., 2021) and long short-term memory networks (LSTMs) (Vu et al., 2020; Yuan et al., 2022). These models utilized

pre-trained word2vec embeddings (Mikolov et al., 2013) and combined neighbouring word representations using convolutional filters or recurrent architectures. Despite their simplicity, some of these models achieved very high-performance baselines that were difficult to surpass with transformer approaches (Ji et al., 2021).

Efforts to apply pre-trained transformers without further modifications to the ICD coding problem were unsuccessful (Ji et al., 2021; Dai et al., 2022). The discharge summary contains 3,594 tokens on average, while the combined set of notes contains an average of 21,916 tokens per patient stay (Ng et al., 2023). Crucial information to predict patient diagnoses is likely to be dispersed throughout these notes, thus models with limited context length risk overlooking a significant portion of relevant data. For this reason, subsequent studies focused on adapting transformer architectures to process longer textual sequences.

PubMedBERT-hier (Ji et al., 2021) employed hierarchical transformers to mitigate the length limitation issue, obtaining substantially better results. This approach segments the document into chunks of 512 tokens and employs a BERT-based model pre-trained on the biomedical domain to encode each segment (Gu et al., 2022). The segments are then combined using a hierarchical transformer running over the CLS-token embeddings. The TrLDC model (Dai et al., 2022) further improved performance by employing a RoBERTa-based model pre-trained from scratch on biomedical articles and clinical notes (Lewis et al., 2020).

The PAAT (Partition Attention) model (Kim, 2022) was able to surpass LSTM-based models like MSMN (Yuan et al., 2022) on the task of identifying the top 50 labels. PAAT combines the Clinical Long-Former and a bi-LSTM, employing partition-based label attention for improved performance. HiLAT (Hierarchical Label-Attention) (Liu et al., 2022) achieved strong results on the top 50 labels by utilizing ClinicalPlusXLNet, which outperforms other transformers like RoBERTa variants, with the downside being that the training speed is four times slower due to its bidirectional context capturing (Liu et al., 2022).

The HTDS (Hierarchical Transformer for Document Sequences) model (Ng et al., 2023) integrated earlier notes into the input context when making decisions about the discharge summary. This model employs a RoBERTa base transformer and a separate transformer layer running over the individual

¹<https://github.com/mireiahernandez/icd-continuous-prediction>

token representations, not only the CLS embeddings. They found that the earlier notes were indeed useful as additional evidence at the end of the hospital stay and provided performance improvements when classifying discharge summaries.

All this prior work has trained systems to assign ICD codes at the end of the hospital stay, whereas we investigate models for making predictions at any point during the stay. Furthermore, while previous work has focused on detecting explicit mentions of diagnoses and treatments in a given text, we investigate to what extent future labels can be inferred based on only earlier documents.

3 Architecture

We investigate a model architecture that can be trained to encode a long temporal sequence of many clinical notes and make predictions at any time point, only using earlier notes as context. The model breaks the sequence into smaller chunks and encodes them using a hierarchical transformer. These chunks are combined with causal label attention, which gathers evidence with label-specific attention heads while ensuring that representations at any time are constructed based on the notes available up to that point, without accessing information in the future. Finally, a probability distribution across the labels is predicted at each possible time point. We refer to the model as a Label-Attentive Hierarchical Sequence Transformer (LAHST) and describe it in more detail below. Figure 1 provides a diagram of the architecture.

Step 1. Document splitting. Each document within the EHR sequence of a patient is tokenized and split into chunks of T tokens. Each patient has a variable total number of chunks, and during training, a maximum of N chunks is selected based on the criteria described in the next section.

Step 2. Chunk encoding. Each of the chunks is encoded with a pre-trained language model (PLM), extracting the CLS-token embedding as the representation, yielding a tensor $e \in \mathbb{R}^{N \times D}$. We use the RoBERTa-base-PM-M3-Voc checkpoint, as it has been trained on two domains that match our task closely: 1) PubMed and PMC, which cover biomedical publications, and 2) MIMIC-III, which contains clinical health records (Lewis et al., 2020).

Step 3. Causal attention. We augment a transformer layer with causal attention (Choromanski et al., 2021) in order to combine temporal information from any previous step without providing

access to information in the future steps. At the same time, the whole sequence can be efficiently processed in parallel by masking any attention connections on the right side of the target position. This component takes as input the sequence of chunk embeddings $e \in \mathbb{R}^{N \times D}$ and generates a sequence of embeddings $h \in \mathbb{R}^{N \times D}$ which combines information over past documents:

$$h_i = \text{CausalAttn}(e_1, \dots, e_i), i \in [1, N] \quad (1)$$

Step 4. Masked multi-head label attention.

We apply label-wise attention (Mullenbach et al., 2018) with two key modifications: the use of multiple attention heads, and the use of causal masking to obtain temporal label-wise document embeddings. For each temporal position t , we define an attention mask $a_t \in \mathbb{R}^{L \times N}$ to prevent attention to future notes, which is constant in the label dimension, and nullifies attention weights beyond temporal position t .

$$a_t[:, i] = \begin{cases} 0, & \text{if } i \leq t \\ -\infty & \text{otherwise} \end{cases} \quad (2)$$

We then combine this mask with multi-head attention (Vaswani et al., 2017) using learnable label embeddings $q \in \mathbb{R}^{L \times D}$ as queries and the previously generated past context embeddings $h \in \mathbb{R}^{N \times D}$ as keys and values:

$$\begin{aligned} d_t &= \text{MultiHeadAttn}(q, h, h, a_t) \\ &= \text{Concat}(\text{head}_1, \dots, \text{head}_H)W^o \end{aligned} \quad (3)$$

Here, each head inputs a linear projection of the key, query and value embeddings $e_{k,i} = W_i^K q$, $e_{q,i} = W_i^Q h$, $e_{v,i} = W_i^V h$ and applies masked attention (Choromanski et al., 2021) as follows:

$$\begin{aligned} \text{head}_i &= \text{Attention}(e_{k,i}, e_{q,i}, e_{v,i}, \text{mask} = a_t) \\ &= \text{SoftMax}\left(\frac{e_{k,i}e_{q,i}^T}{\sqrt{D/H}} + a_t\right)e_{v,i} \end{aligned} \quad (4)$$

This yields a sequence of label-wise document embeddings $d_t \in \mathbb{R}^{L \times D}$ for each position $t \in \{1, \dots, N\}$. In practice, we obtain all the embeddings $d \in \mathbb{R}^{N \times L \times D}$ efficiently in one pass by assigning the batch dimension to the temporal dimension.

Step 5. Temporal label-wise predictions. Finally, temporal probabilities are calculated by projecting the embedding using linear weights $w \in$

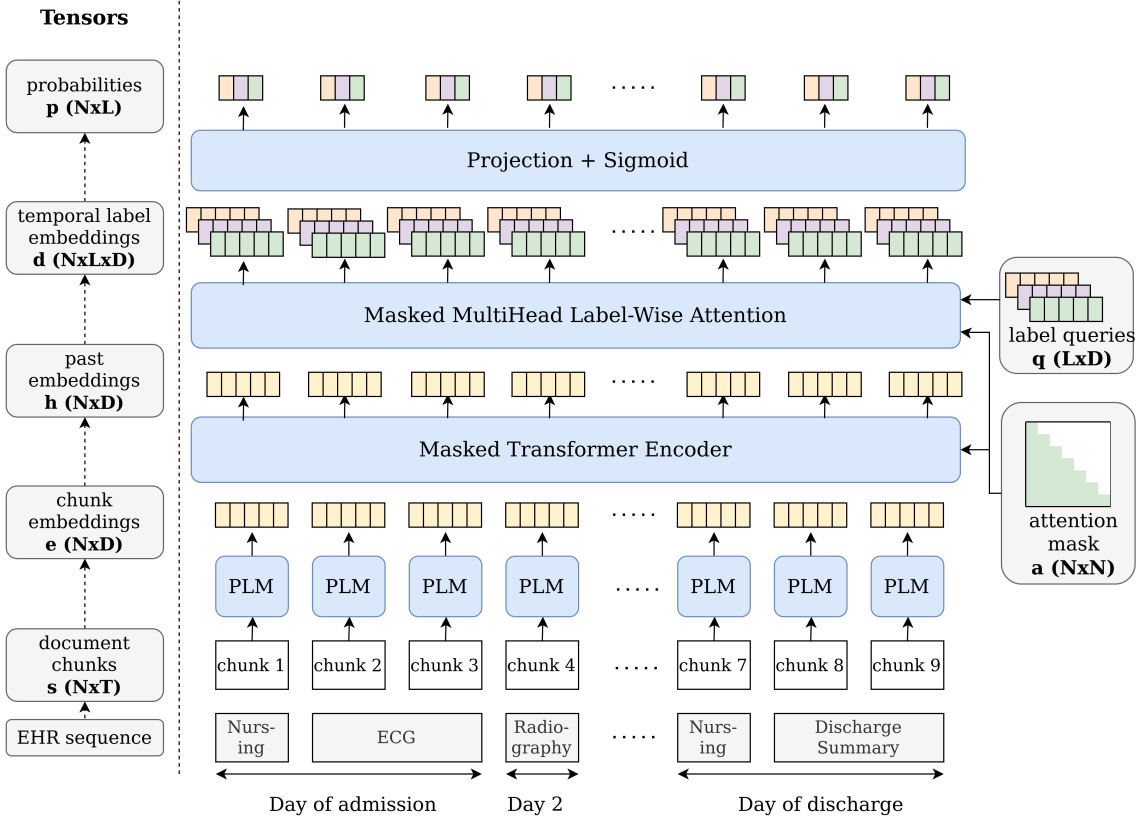


Figure 1: **LAHST** (Label-Attentive Hierarchical Sequence Transformer) architecture. Clinical notes generated throughout the hospital stay are split into chunks. Each chunk is encoded using a pre-trained language model (PLM) to extract the CLS-token embedding. Next, a hierarchical transformer encoder is applied, utilizing causal masking to combine information among past segment embeddings. Finally, the network generates a distinct document representation for each label and temporal point combination and these are then transformed into probabilities by the output layer.

$\mathbb{R}^{L \times D}$ followed by a sigmoid activation. The probability at time t for label l is calculated using the the label weight $w_l \in \mathbb{R}^D$ and the label document embedding at position t , denoted as $d_{t,l} \in \mathbb{R}^D$:

$$p_{t,l} = \text{sigmoid}(w_l \cdot d_{t,l}) \quad (5)$$

The output of the model is a probability matrix $p \in \mathbb{R}^{N \times L}$, containing probabilities for each label at each temporal point. The masking process within the transformer and label attention modules ensures that time t probability calculations consider only past documents. The model is trained using the binary cross-entropy loss.

4 Extending the Context

Hierarchical transformer architectures break long inputs into smaller components and reduce the number of long-distance attention operations, thereby keeping memory and computation requirements more manageable when processing very long sequences. This makes them well-suited for ICD

code classification, as local context is more important for this task and hierarchical models have been shown to outperform long-context models in this setting (Dai et al., 2022). However, even hierarchical models have difficulty with very long sequences, particularly during training. The gradient must be backpropagated through each individual chunk encoding, which can easily cause memory issues when the models are large and the number of segments exceeds a maximum limit.

For this reason, we propose a novel solution for applying hierarchical transformers to very long document sequences, such as the sequences of notes in EHR. We refer to this method as the Extended Context Algorithm (ECA). It consists of the following modifications to the training and inference loops.

Training (Algorithm 1). For each episode of training, the loop iterates over the training dataset \mathcal{D}_{train} , processing each data sample (s, y) , where s is the input sequence and y is the corresponding label. Within the loop, a random selection of

Algorithm 1 ECA Training loop

$\mathcal{D}_{train} \leftarrow$ training set (sequence-label pairs)
 $N_{max} \leftarrow$ max. number of chunks
for each episode **do**
 for each (s, y) in \mathcal{D}_{train} **do**
 $m \leftarrow \min(N_{max}, \text{len}(s))$
 select m random indices i_1, \dots, i_m
 sort i_1, \dots, i_m in ascending order
 $s' \leftarrow [s[i_1], \dots, s[i_m]]$
 $p, h \leftarrow \text{model.forward}(s')$
 $\mathcal{L} \leftarrow BCE(y, p)$
 do backward pass and optimizer step
 end for
end for

Algorithm 2 ECA Inference loop

$\mathcal{D}_{test} \leftarrow$ test set (no labels)
 $N_{max} \leftarrow$ max. number of chunks
for each s in \mathcal{D}_{test} **do**
 $h_{list} \leftarrow$ empty list
 for i in range(0, $\text{len}(s)$, N_{max}) **do**
 $s_{batch} \leftarrow s[i : i + N_{max}]$
 $p_{batch}, h_{batch} \leftarrow \text{model.forward}(s_{batch})$
 append h_{batch} to h_{list}
 end for
 $h \leftarrow$ concatenate h_{list} along batch dim.
 $p \leftarrow \text{model.label_attention}(h)$
end for

notes is chosen to create a subset of the input sequence, with the maximum number of chunks set as N_{max} . These sub-sequences s' are then used for optimizing the model, each time sampling a slightly different training instance. Instead of trying to fit the whole sequence into the input during training, we sample notes and form multiple different shorter versions of the sequence for training the model. This has the added benefit of creating a data augmentation effect, as the model learns to make decisions based on different versions of the same datapoint.

Inference (Algorithm 2). During inference, we process all the notes in the sequence in batches of N_{max} chunks. Each sequence batch, denoted as s_{batch} , is encoded to obtain embeddings $h_{batch} \in \mathbb{R}^{N_{max} \times D}$. Even if the full sequence does not fit into memory, it can be processed in separate batches to obtain all the h_{batch} embeddings. These embeddings are then concatenated along the batch dimension to obtain chunk embeddings for the complete sequence $h \in \mathbb{R}^{N_{total} \times D}$. Finally, the

collected embeddings are passed through causal attention and masked multi-head label attention to obtain predictions p based on the complete sequence.

As the computation can be performed in separate batches and then combined, this allows for considerably longer sequences to be used as input during inference. Unlike other methods for extending the context of transformers that rely on reducing or compressing long-distance attention (Beltagy et al., 2020; Munkhdalai et al., 2024), this proposed method is also exact – the result is always the same as it would be with a single pass using infinite memory.

5 Experiment Set-up

5.1 Evaluation framework

We investigate the novel task of temporal ICD code prediction, which requires the prediction of ICD codes at any point during the hospital stay using the notes available at that time, without relying on the discharge summary. To evaluate the performance, we will compare the predictive power of our model at different points throughout the EHR sequence. Our evaluation setup is inspired by the ClinicalBERT model (Huang et al., 2019), which evaluates the likelihood of readmission at different cut-off times since admission.

The cut-off times were selected to be the 25%, 50%, and 75% percentiles of the total volume of notes present in the training dataset, which are shown in Table 2. These correspond to 2, 5, and 13-day cut-offs, respectively. For example, in the 2-day setting, the model only has access to the notes written in the first 2 days in order to predict all the ICD codes that will be assigned to that patient by the end of their hospital stay. Where space allows, we additionally report on all the notes up to (but excluding) the discharge summary, indicating a setting where the model could be used to assist in the writing of the discharge summary itself. For comparison, we also report performance on the full sequence which includes the discharge summary, although this setting is retrospective and would not provide any predictive benefit. In line with widely used approaches to ICD coding (Mullenbach et al., 2018), we focus on Micro-F1, Micro-AUC and Precision@5 metrics, with additional metrics provided in the appendix.

Percentile	Days elapsed	# notes
25%	1.8	112,594
50%	5.2	225,160
75%	12.8	337,726

Table 2: Percentiles of the total volume of notes present in the training dataset. The number of days corresponding to the 25th, 50th, and 75th percentiles will be used as temporal evaluation points throughout this project.

	# chunks / patient	# patients
2 days	17.9 \pm 22.1	1,559
5 days	27.6 \pm 33.9	1,559
13 days	35.8 \pm 42.4	1,559
excl. DS	40.4 \pm 47.7	1,559
last day	48.4 \pm 48.1	1,573

Table 3: Length of EHR in number of chunks per patient (average and standard deviation) and count of patients of our dataset at different temporal cut-offs (dev set).

5.2 Preprocessing

We use the MIMIC-III dataset (Johnson et al., 2016) for evaluation, as it contains a collection of Electronic Health Records with timestamped free-text reports by nurses and doctors, together with the corresponding ICD-9 labels. First, we follow the preprocessing steps outlined by the CAML approach (Mullenbach et al., 2018) to obtain a dataset of free-text clinical notes paired with ICD diagnoses and procedure codes, and we also extract their proposed train/dev/test splits. The label space is vast, so following their method, we focus on predicting the top 50 codes.

For our novel task, we perform some additional preprocessing steps. First, we extract the timestamps of each note and, in cases where the specific time is missing, assign it to 12:00:00 of that day. Moreover, we found that some patients had additional notes beyond the discharge summary document, such as other discharge summaries or nursing notes. We exclude these additional notes to ensure that our EHR sequence always concludes with a single discharge summary document. We also exclude 14 patients as their EHR contains no other notes besides the discharge summary. Table 3 displays the statistics of our dataset at various temporal cut-offs.

5.3 Implementation details

The model is implemented in Pytorch and it was trained on an Nvidia GeForce GTX Titan Xp (12GB RAM) GPU, utilizing an average memory of 11.22 GB. The model processed 5 samples per second and training took an average time of 11 hours and 50 minutes. We used a super-convergence learning rate scheduler (Smith and Topin, 2019), based on its use in HTDS (Ng et al., 2023), and an early-stopping strategy with a 3 epoch patience and a maximum of 20 epochs. Chunk size T was set to 512 tokens as that is the largest size supported by RoBERTa-base-PM-M3-Voc. For the main experiments, a limit of $N_{max} = 16$ was used during training, while the entire sequence (with up to 181 chunks) was used for inference. The tuning ranges and chosen hyperparameter values are included in Appendix A.

6 Results

In addition to the LAHST framework described in Sections 3 and 4, we also evaluate PubMedBERT-Hier (Ji et al., 2021) and HTDS (Ng et al., 2023) on the early prediction task. HTDS was trained to consider earlier notes in the context while making decisions about the discharge summary, making it the most likely existing model to also perform well on the early prediction task. In addition, HTDS results are very close to the state-of-the-art on the MIMIC-III dataset, making it a very strong baseline. However, HTDS is a larger model and requires considerably more GPU resources compared to LAHST. Therefore, we also report a modified version (HTDS*) which has a comparable number of parameters. We also include the performance of TrLDC (Dai et al., 2022) from the respective paper as an additional strong baseline on classification of the discharge reports.

In Table 4 we report the performance of these systems at increasingly challenging temporal cut-offs. LAHST shows strong performance at any time point, outperforming all the other models at every early prediction task. The results indicate that some of the diagnosis and treatment codes for the whole hospital stay can be predicted already within the first few days of admission. While the performance of all systems is expectedly lower in the more challenging settings, they are still able to reach 46% F1 and 82.9% AUC with only 2 days of information, which could provide useful pre-

	Last day			0-13 days			0-5 days			0-2 days		
Model	F1	AUC	P@5	F1	AUC	P@5	F1	AUC	P@5	F1	AUC	P@5
TrLDC	70.1	93.7	65.9	-	-	-	-	-	-	-	-	-
PMB-H	67.2	91.5	63.0	30.7	68.0	30.2	31.3	68.4	31.0	31.7	68.7	31.5
HTDS	73.3	95.2	68.1	49.7	82.1	47.6	47.5	80.6	45.9	44.5	78.7	43.6
HTDS*	70.7	93.8	66.2	48.6	82.0	47.0	46.7	80.7	45.5	43.6	78.7	43.3
LAHST	70.3	94.6	67.5	52.9	87.0	53.0	50.3	85.4	50.7	46.0	82.9	47.1

Table 4: Evaluation on the early ICD code prediction task at increasingly challenging temporal cut-offs. TrLDC (Dai et al., 2022) result is from the respective paper. We evaluated PubMedBERT-Hier (PMB-H; Ji et al., 2021) and HTDS (Ng et al., 2023) at different early prediction points. HTDS* is a version of HTDS that is more comparable to LAHST in terms of computation requirements. LAHST is the model described in Sections 3 and 4. Results for PMB-H, HTDS, HTDS* and LAHST are averaged over 3 runs with different random seeds.

dictions to the hospital staff. In all the early prediction settings, LAHST achieves the best results according to all metrics. While HTDS is trained to look at earlier documents and is also able to make competitive predictions, it is reliant on information in the discharge summary and therefore underperforms when this is not available. In contrast, the LAHST model is trained to make predictions based on varying amounts of evidence and achieves the best performance.

In the "Last day" setting, which includes the discharge summary, HTDS slightly outperforms LAHST – this is expected, as HTDS is a larger model and specifically trained for discharge summaries. However, when compared to the similarly-sized HTDS*, LAHST delivers comparable F1 along with improved AUC and P@5. Even though LAHST is not trained for this particular setting, the supervision on earlier time points helps it achieve good results also when classifying discharge summaries. In addition, it outperforms both PubMedBERT-Hier and TrLDC according to all metrics. We include larger results tables with additional metrics in Appendix B.

7 Analyzing the Extended Context

Selection of context during inference. The Extended Context Algorithm (ECA) allows the model to include much longer EHR sequences in the context during inference (with a generous 181 chunk cap applied in our experiments). We evaluate the effect of this algorithm compared to alternative strategies used in other hierarchical models. The "Last" setting uses the most recent 16 chunks of text, illustrating the setting where the sequence is truncated from the beginning in order to fit into the

	Last	Random	ECA
<i>0-2 days</i>	33.5 \pm 0.6	29.2 \pm 0.5	46.2 \pm 0.1
<i>0-5 days</i>	37.6 \pm 0.3	35.1 \pm 0.4	50.9 \pm 0.1
<i>0-13 days</i>	38.2 \pm 0.1	37.7 \pm 0.5	53.6 \pm 0.2
<i>Excl.DS</i>	37.9 \pm 0.2	38.4 \pm 0.4	54.3 \pm 0.2
<i>Last day</i>	71.0 \pm 0.3	71.3 \pm 0.2	71.1 \pm 0.1

Table 5: Micro-F1 score on the development set, using the LAHST model with alternative strategies for context inclusion.

model. The "Random" setting samples a random subset of chunks from the sequence instead. The results in Table 5 show that processing the entire sequence with ECA yields substantial performance improvements (+12.7, +13.3 and +15.4 Micro-F1 score for 2 days, 5 days and 13 days) compared to truncating or sampling the sequence. This result highlights the importance of including all the available notes in the input. Only when the discharge note is available (in the 'Last day' setting) the previous notes become less important and all the strategies give the same performance.

Selection of context during training. We investigate the effect of randomly sampling different sub-sequences of notes during training. We train an alternative version of LAHST by truncating the sequence to the most recent notes instead of sampling them randomly. During inference, both versions still receive all the notes as input, as described in Algorithm 2. The results in Table 6 show how training without random sampling substantially decreases performance across all evaluation points (-8.4, -8.4, -8.2, -8.0, -0.9, F1 respectively). This indicates that randomly sampling different sub-sequences during optimization augments the training data with

	Last	Random/ECA
<i>0-2 days</i>	37.8 \pm 0.2	46.2 \pm 0.1
<i>0-5 days</i>	42.5 \pm 0.1	50.9 \pm 0.1
<i>0-13 days</i>	45.4 \pm 0.1	53.6 \pm 0.2
<i>Excl. DS</i>	46.3 \pm 0.1	54.3 \pm 0.2
<i>Last day</i>	70.2 \pm 0.2	71.1 \pm 0.1

Table 6: Micro-F1 of LAHST on the development set, using alternative sampling strategies during training.

different variations which helps the model better generalize to different temporal cut-offs, without increasing memory or computation requirements.

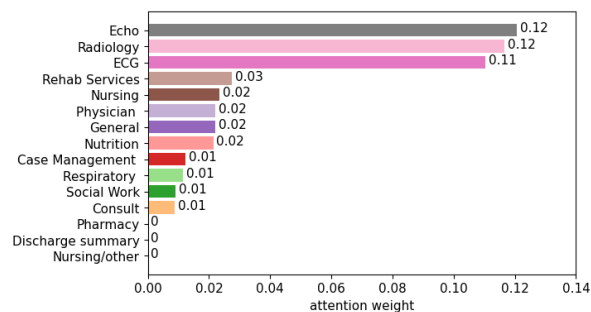
8 Model Interpretability

The attention weights in the label-attention layer of LAHST can potentially be used as an importance indicator of different input notes. A higher weight is associated with an increased relevance of the particular document to predicting a specific code. For an initial visualization, we average the weights across all the codes to find which document types are most important at different temporal cut-offs.

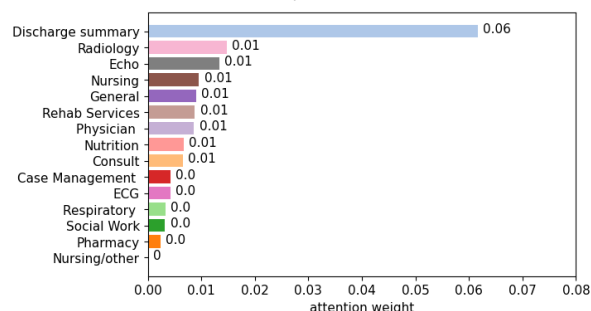
The results are shown in Fig. 2. Within the 2-day cut-off, all the reports that have diagnostic characteristics have received the highest attention weights. For example, the echocardiography report is the description of an ultrasound test to identify abnormalities in the heart structure and is used by cardiologists to diagnose heart diseases (Van et al., 2023). The radiology reports detail the results of imaging procedures such as X-rays and MRIs to diagnose diseases (Alarifi et al., 2021). All of such reports are highly technical and are specifically created to assist physicians with diagnostic practices. In the absence of the discharge summary, they are the most valuable document types for making early predictions of ICD-9 codes and the network has correctly focused more attention on them. In the "Last day" setting, the discharge summary becomes available, containing an overview of the entire hospital stay, and the same model is able to switch most of its attention to it.

9 Conclusions

In this study, we investigated the potential of predicting ICD codes for the whole patient stay at different time points during their stay. Being able to predict likely diagnoses and treatments in advance would have important applications for predictive medicine, by enabling early diagnosis, sugges-



(a) 2-days cut-off



(b) Last day cut-off

Figure 2: Average attention weight per document type at different temporal cut-offs. The LAHST model processes the complete EHR sequence and focuses more on reports of diagnostic tests for early prediction, switching to the discharge summary when it is available.

tions for treatments, and optimization of resource allocation. We designed a specialized architecture (LAHST) for this task, which uses a hierarchical structure combined with label attention and causal attention to efficiently make predictions at any possible time points in the EHR sequence. The Extended Context Algorithm was further proposed to allow the model to better handle very long sequences of notes. The system is trained by sampling different sub-sequences of notes, which allows the model to fit into memory while also augmenting the data with variations of available examples. During inference, the whole sequence is then processed separately in batches and combined together with a single attention layer, allowing for lossless representations of very long context to be calculated.

Our experiments showed that useful predictions regarding the final ICD codes for a patient can be made already soon after the hospital admission. The LAHST model substantially outperformed existing approaches on the early prediction task, while also achieving competitive results on the standard task of assigning codes to discharge summaries. The model achieved 82.9% AUC already 2 days after admission, indicating that it is

able to rank and suggest relevant ICD codes based on limited information very early into a hospital stay.

10 Limitations

The primary focus of this project was to investigate the feasibility of this novel task and explore a novel architecture for the early prediction of ICD codes. Even though this could open up new avenues for early disease detection and procedure forecasting, our work has some limitations that should be considered in future work.

Firstly, our study is limited to the MIMIC dataset as it is one of the largest and most established available datasets containing electronic health records and ICD codes. However, the findings based on this dataset may not generalise equally to every clinical setting. Therefore, new experiments would need to be conducted on representative data samples before considering applying such technology in practice.

Our experiments focused on PubMedBERT-Hier (Ji et al., 2021), HTDS (Ng et al., 2023) and LAHST. However, there are many other architectures and pre-trained models available which could be investigated in this setting.

Our model is based on a hierarchical transformer architecture which achieves good performance but is also quite computationally expensive compared to LSTM or CNN-based approaches (training the model took roughly 12 hours on a 12GB GPU). With our computational resources, we were limited to running experiments using the [CLS]-token representation and a maximum of 16 chunks in a batch. However, with additional resources this work could be further scaled up by retaining all token representations and increasing the model size to allow for the allocation of additional chunks.

Finally, our evaluation of the temporal ICD coding task is focused on reporting the aggregate metrics for the top 50 ICD-9 coding labels. Future work could investigate a larger number of labels, along with analysing the performance separately on individual labels and label types.

11 Ethics Statement

After careful consideration, we have determined that no ethical conflicts apply to this project. While clinical data is inherently sensitive, it is important to note that the MIMIC-III dataset has undergone a rigorous de-identification process, following the guidelines outlined by the Health Insurance Porta-

bility and Accountability Act (HIPAA). This de-identification process ensures that the dataset can be used for research purposes on an international scale (Johnson et al., 2016).

While no conflicts were identified, machine learning systems for ICD coding carry certain risks when deployed in hospitals. Firstly, automated approaches are trained in a supervised manner using data from hospitals, and therefore, they are susceptible to reproducing manual coding errors. These errors may include miscoding due to misunderstandings of abbreviations and synonyms or overbilling due to unbundling errors (Sonabend W et al., 2020). Moreover, automated systems may also suffer from distribution shifts, potentially affecting their portability across various EHR systems in different hospitals (Sonabend W et al., 2020). To address these concerns, it is important to build interpretable models and develop tools that enable human coders to supervise the decisions made by ICD coding models.

Acknowledgements

Mireia Hernandez Caralt acknowledges that the project that gave rise to these results received the support of a fellowship from “la Caixa” Foundation (ID 100010434). The fellowship code is LCF/BQ/EU22/11930076.

References

- Mohammad Alarifi, Patrick Timothy, Jabour Abdulrahman, Min Wu, and Jake Luo. 2021. [Understanding patient needs and gaps in radiology reports through online discussion forum analysis](#). *Insights Imaging*, 12(50).
- Iz Beltagy, Matthew E. Peters, and Arman Cohan. 2020. [Longformer: The long-document transformer](#). *CoRR*, abs/2004.05150.
- Donna J Cartwright. 2013. Icd-9-cm to icd-10-cm codes: What? why? how? *Advances in Wound Care*, 2(10):588–592.
- Hua Cheng, Rana Jafari, April Russell, Russell Klopfer, Edmond Lu, Benjamin Striner, and Matthew Gormley. 2023. [MDACE: MIMIC documents annotated with code evidence](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7534–7550, Toronto, Canada. Association for Computational Linguistics.
- Krzysztof Choromanski, Han Lin, Haoxian Chen, Tianyi Zhang, Arijit Sehanobish, Valerii Likhoshervstov, Jack Parker-Holder, Tamás Sarlós, Adrian

- Weller, and Thomas Weingarten. 2021. [From block-toeplitz matrices to differential equations on graphs: towards a general theory for scalable masked transformers](#). In *International Conference on Machine Learning*.
- Xiang Dai, Ilias Chalkidis, Sune Darkner, and Desmond Elliott. 2022. [Revisiting transformer-based models for long document classification](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 7212–7230, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Yu Gu, Robert Tinn, Hao Cheng, Michael Lucas, Naoto Usuyama, Xiaodong Liu, Tristan Naumann, Jianfeng Gao, and Hoifung Poon. 2022. [Domain-Specific Language Model Pretraining for Biomedical Natural Language Processing](#). *ACM transactions on computing for healthcare*, 3(1):1–23.
- Kexin Huang, Jaan Alntosaar, and Rajesh Ranganath. 2019. [ClinicalBERT: modeling clinical notes and predicting hospital readmission](#). *Computing Research Repository*, arXiv:1904.05342v3. Version 3.
- Peter B Jensen, Lars J Jensen, and Søren Brunak. 2012. Mining electronic health records: towards better research applications and clinical care. *Nature Reviews Genetics*, 13(6):395–405.
- Shaoxiong Ji, Matti Hölttä, and Pekka Marttinen. 2021. [Does the magic of bert apply to medical code assignment? a quantitative study](#). *Computers in Biology and Medicine*, 139:104998.
- Alistair E. W. Johnson, Tom J. Pollard, Lu Shen, Li-Wei H. H. Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G. Mark. 2016. [MIMIC-III, a freely accessible critical care database](#). *Scientific Data*, 3(160035).
- Haanju Yoo Daeseong Kim Sewon Kim. 2022. [An Automatic ICD Coding Network Using Partition-Based Label Attention](#). *SSRN Electronic Journal*.
- Patrick Lewis, Myle Ott, Jingfei Du, and Veselin Stoyanov. 2020. [Pretrained language models for biomedical and clinical tasks: understanding and extending the state-of-the-art](#). *Proceedings of the 3rd Clinical Natural Language Processing Workshop*, page 146.
- Fei Li and Hong Yu. 2020. [ICD Coding from Clinical Text Using Multi-Filter Residual Convolutional Neural Network](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):8180.
- Leibo Liu, Oscar Perez-Concha, Anthony Nguyen, Vicki Bennett, and Louisa Jorm. 2022. [Hierarchical label-wise attention transformer model for explainable icd coding](#). *Journal of biomedical informatics*, 133:104161.
- Yang Liu, Hua Cheng, Russell Klopfer, Matthew R. Gormley, and Thomas Schaaf. 2021. [Effective convolutional attention network for multi-label clinical document classification](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 5941–5953, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Tomas Mikolov, Kai Chen, G.s Corrado, and Jeffrey Dean. 2013. [Efficient estimation of word representations in vector space](#). *Proceedings of Workshop at ICLR*, 2013.
- James Mullenbach, Sarah Wiegrefe, Jon Duke, Jimeng Sun, and Jacob Eisenstein. 2018. [Explainable prediction of medical codes from clinical text](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1101–1111, New Orleans, Louisiana. Association for Computational Linguistics.
- Tsendsuren Munkhdalai, Manaal Faruqui, and Siddharth Gopal. 2024. [Leave no context behind: Efficient infinite context transformers with infinite attention](#). *arXiv preprint arXiv:2404.07143*.
- Boon Liang Clarence Ng, Diogo Santos, and Marek Rei. 2023. [Modelling temporal document sequences for clinical ICD coding](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 1640–1649, Dubrovnik, Croatia. Association for Computational Linguistics.
- Leslie N Smith and Nicholay Topin. 2019. [Very fast training of neural networks using large learning rate](#). *Artificial intelligence and machine learning for multi-domain operations applications*, 1106:369–386.
- Aaron Sonabend W, Winston Cai, Yuri Ahuja, Ashwin Ananthakrishnan, Zongqi Xia, Sheng Yu, and Chuan Hong. 2020. [Automated ICD coding via unsupervised knowledge integration \(UNITE\)](#). *International journal of medical informatics*, 139, 104135.
- Maryam Tayefi, Phuong Ngo, Taridzo Chomutare, Hercules Dalianis, Elisa Salvi, Andrius Budrionis, and Fred Godtliebsen. 2021. [Challenges and opportunities beyond structured data in analysis of electronic health records](#). *Wiley Interdisciplinary Reviews: Computational Statistics*, 13(6):e1549.
- Phi Nguyen Van, Hieu Pham Huy, and Long Tran Quoc. 2023. [Echocardiography segmentation using neural ode-based diffeomorphic registration field](#). *IEEE Transactions On Medical Imaging*, XX.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems (NIPS)*.
- Thanh Vu, Dat Nguyen, and Anthony Nguyen. 2020. [A label attention model for ICD coding from clinical text](#). In *Proceedings of IJCAI*. Doi:10.24963/ijcai.2020/461.

Zheng Yuan, Chuanqi Tan, and Songfang Huang. 2022. Code synonyms do matter: Multiple synonyms matching network for automatic ICD coding. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 808–814, Dublin, Ireland. Association for Computational Linguistics.

Appendix A

Hyper-parameter	Range
Num. Layers (Mask. Transf.)	1,2,3
Num. Heads (Mask. Transf.)	1,2,3
Num. Heads (Label Atten.)	1,2,3
Peak LR	1e-5, 5e-5 , 1e-4

Table 7: The range of hyperparameters searched for tuning the model. The chosen value is shown in bold.

Appendix B

Detailed results tables using different time cut-offs.

	0-2 days				
	Micro-F1	Macro-F1	Micro-AUC	Macro-AUC	P@5
PubMedBERT-Hier (Ji et al., 2021)	-	-	-	-	-
TrLDC (Dai et al., 2022)	-	-	-	-	-
HTDS (Ng et al., 2023)	44.5	39.7	78.7	77.5	43.6
HTDS*	43.6	40.0	78.7	76.1	43.3
LAHST	46.0	40.1	82.9	79.5	47.1

	0-5 days				
	Micro-F1	Macro-F1	Micro-AUC	Macro-AUC	P@5
PubMedBERT-Hier (Ji et al., 2021)	-	-	-	-	-
TrLDC (Dai et al., 2022)	-	-	-	-	-
HTDS (Ng et al., 2023)	47.5	42.5	80.6	79.5	45.9
HTDS*	46.7	42.5	80.7	78.2	45.5
LAHST	50.3	44.6	85.4	82.2	50.7

	0-13 days				
	Micro-F1	Macro-F1	Micro-AUC	Macro-AUC	P@5
PubMedBERT-Hier (Ji et al., 2021)	-	-	-	-	-
TrLDC (Dai et al., 2022)	-	-	-	-	-
HTDS (Ng et al., 2023)	49.7	44.6	82.1	81.2	47.6
HTDS*	48.6	44.6	82.0	79.7	47.0
LAHST	52.9	47.3	87.0	83.8	53.0

	Excl DS				
	Micro-F1	Macro-F1	Micro-AUC	Macro-AUC	P@5
PubMedBERT-Hier (Ji et al., 2021)	-	-	-	-	-
TrLDC (Dai et al., 2022)	-	-	-	-	-
HTDS (Ng et al., 2023)	50.2	45.2	82.3	81.5	47.8
HTDS*	49.0	45.1	82.4	80.2	47.5
LAHST	53.5	47.8	87.3	84.1	53.7

	Last Day				
	Micro-F1	Macro-F1	Micro-AUC	Macro-AUC	P@5
PubMedBERT-Hier (Ji et al., 2021)	68.1	63.3	90.8	88.6	64.4
TrLDC (Dai et al., 2022)	70.1	63.8	93.7	91.4	65.9
HTDS (Ng et al., 2023)	73.3	67.7	95.2	93.6	68.1
HTDS*	70.7	64.9	93.8	91.6	66.2
LAHST	70.3	64.3	94.6	92.6	67.5