

Accurate and Well-Calibrated ICD Code Assignment Through Attention Over Diverse Label Embeddings

Gonçalo Gomes Isabel Coutinho Bruno Martins

Instituto Superior Técnico and INESC-ID, University of Lisbon

{goncaloecgomes, isabel.coutinho, bruno.g.martins}@tecnico.ulisboa.pt

Abstract

Although the International Classification of Diseases (ICD) has been adopted worldwide, manually assigning ICD codes to clinical text is time-consuming, error-prone, and expensive, motivating the development of automated approaches. This paper describes a novel approach for automated ICD coding, combining several ideas from previous related work. We specifically employ a strong Transformer-based model as a text encoder and, to handle lengthy clinical narratives, we explored either (a) adapting the base encoder model into a Longformer, or (b) dividing the text into chunks and processing each chunk independently. The representations produced by the encoder are combined with a label embedding mechanism that explores diverse ICD code synonyms. Experiments with different splits of the MIMIC-III dataset show that the proposed approach outperforms the current state-of-the-art models in ICD coding, with the label embeddings significantly contributing to the good performance. Our approach also leads to properly calibrated classification results, which can effectively inform downstream tasks such as quantification.

1 Introduction

The International Classification of Diseases (ICD¹) coding system, proposed by the World Health Organization, stands as a universal standard for precise documentation in the medical domain (O'Malley et al., 2005). Still, the manual assignment of ICD codes to clinical text is a time-consuming, labor intensive, and error-prone task, which has led to the exploration of automated methods, e.g. using deep learning algorithms for text classification.

Despite many previous efforts, automatic ICD coding is still challenging. Clinical notes consist of long text narratives that use a specialized medical

vocabulary and are associated with a high dimensional, sparse, and imbalanced label space.

In addition to accurately classifying individual clinical notes, estimating the prevalence of ICD codes within a dataset is also important for many practical applications. This corresponds to a text quantification problem (Schumacher et al., 2021; Moreo et al., 2022), for which access to properly calibrated text classification models can be helpful.

This paper describes a novel approach to ICD coding, combining several ideas from previous work. In particular, we explored two different text encoding strategies using a strong Transformer-based model (Yang et al., 2022a), explicitly dealing with the lengthy nature of documents like hospital discharge summaries. The resulting representations are combined with a label embedding mechanism inspired by the proposal from Yuan et al. (2022), which explores diverse ICD code synonyms. Additionally, taking inspiration from the MLP-based quantification approach from Coutinho and Martins (2023), we present a training setup in which multi-label classification and text quantification are jointly addressed. This additional step aims to improve model calibration while also informing downstream tasks such as text quantification.

Following previous studies, the proposed model was evaluated on the publicly available MIMIC-III dataset (Johnson et al., 2016), specifically analyzing results on two subsets of hospital discharge summaries, namely MIMIC-III-50 (Mullenbach et al., 2018) and MIMIC-III-clean (Edin et al., 2023). Our approach surpasses standard baselines and previous state-of-the-art models for ICD coding across all evaluated metrics, while simultaneously providing interesting results regarding model calibration and ICD code quantification. The source code supporting our experiments is available in a GitHub repository².

¹<https://www.who.int/standards/classifications/classification-of-diseases>

²https://github.com/gecgomes/ICD_Coding_MSAM

The remaining parts of this paper are organized as follows: Section 2 reviews the related literature, while Section 3 introduces our novel approach for ICD coding and quantification. Section 4 presents the experimental results, establishing a direct comparison with previous studies. Finally, Section 5 summarizes our contributions and discusses future research directions. The paper ends with a discussion on limitations and ethical considerations.

2 Related Work

Several previous studies have addressed the problem of automatic ICD coding. For instance, [Mullenbach et al. \(2018\)](#) introduced the Convolutional Attention for Multi-Label classification (CAML) approach, which is still commonly considered as a baseline. CAML employs a label-wise attention mechanism that enables the model to learn distinct document representations for each label, selecting relevant parts of the document for each ICD code. The authors conducted experiments on MIMIC datasets ([Lee et al., 2011](#); [Johnson et al., 2016](#)), and the data splits developed for this work were made publicly available. This study is considered an essential milestone for reproducibility.

Aiming to address CAML’s limitations in capturing variable-sized text patterns, [Xie et al. \(2019\)](#) improved the convolutional attention model by introducing a densely connected CNN with multi-scale feature attention (MSATT-KG), which produces variable n -gram features and adaptively selects informative features based on neighborhood context. This method also incorporates a graph CNN to capture hierarchical relationships among medical codes. In turn, [Li and Yu \(2020\)](#) proposed MultiResCNN, i.e. a tailored CNN architecture combining multi-filter convolutions and residual convolutions, capturing patterns of different lengths and achieving superior performance over CAML.

[Vu et al. \(2020\)](#) introduced LAAT, i.e., a model that combines an RNN-based encoder with a new label attention mechanism for ICD coding. LAAT aimed to handle the variability in text segment lengths and the interdependence among different segments related to ICD codes. Additionally, the authors introduced a hierarchical joint learning mechanism to address the class imbalance issue.

[Yuan et al. \(2022\)](#) put forth the Multiple Synonyms Matching Network (MSMN) as an alternative approach to ICD coding. Rather than relying on the ICD code hierarchy, the authors leveraged

synonyms to enhance code representation learning and improve coding performance.

Nowadays, Transformer-based Language Models (LMs) are becoming the fundamental technology for medical AI systems to process clinical narratives. For instance, [Yang et al. \(2022a\)](#) developed a large clinical LM named GatorTron, training a Transformer encoder model on text narratives from de-identified clinical notes from University of Florida Health, PubMed articles, and Wikipedia. The authors also examined how increasing the number of parameters enhances the performance on various Natural Language Processing (NLP) tasks, including named entity recognition, medical relation extraction, semantic textual similarity, natural language inference, and medical question answering. Their results showed that GatorTron outperformed previous Transformer models across various NLP tasks within the biomedical and clinical domains.

[Dai et al. \(2022\)](#) compared Transformer models for long document classification, focusing on mitigating the computational overheads associated with encoding large texts. In turn, [Huang et al. \(2022\)](#) investigated limitations associated with using pre-trained Transformer-based LMs, identifying challenges regarding large label spaces, long input lengths, and domain disparities. The authors proposed PLM-ICD, i.e., a framework that effectively handles these challenges and achieves superior results on the MIMIC-III dataset, surpassing previously existing methods.

In a recent study, [Edin et al. \(2023\)](#) argued that the proper assessment of model performance on ICD coding had often struggled with weak configurations, poorly designed train-test splits, and inadequate evaluation procedures. The authors pinpointed significant issues with the MIMIC-III splits released by [Mullenbach et al. \(2018\)](#). They proposed a new dataset split using stratified sampling to ensure a complete representation of all classes, referred to as MIMIC-III-clean.

Regarding text quantification, various algorithms have been proposed in recent years ([Schumacher et al., 2021](#)). Still, few previous studies have specifically considered multi-label settings ([Moreo et al., 2022](#)). [Coutinho and Martins \(2023\)](#) explored the use of a Multi-Layer Perceptron (MLP) model for ICD code quantification, taking inspiration from under-complete denoising auto-encoders. The MLP was trained to refine estimates provided by the Probabilistic Classify and Count (PCC) method, considering label correlations. Experiments with

different MIMIC-III dataset splits showed that the proposed method outperforms baseline approaches such as Classify and Count (CC) and PCC.

3 Proposed Approach

This work presents a novel approach for ICD coding, aiming at strong classification performance together with well-calibrated outputs, which can inform downstream tasks such as text quantification.

3.1 Clinical Text Modeling

Within the proposed approach, we compared two different text encoding strategies to handle long clinical documents, using GatorTron-base as the foundational Language Model (LM) for both strategies. GatorTron-base is a Megatron BERT model pre-trained on the healthcare domain, previously described by Yang et al. (2022a). This model is publicly available in the NVIDIA NGC³ Catalog and also through the HuggingFace⁴ library.

As a first strategy, we considered a Longformer Encoding (LE) approach, where the standard self-attention mechanism of GatorTron-base is replaced by the Longformer sparse attention mechanism, which scales linearly with sequence length (Beltagy et al., 2020). This makes it possible to process documents of thousands of tokens. To adapt the GatorTron-base into a Longformer, we resized the positional embeddings to match the new maximum length allowed for the input sequences (in our case, equal to 8,192 tokens). The LE uses a local windowed self-attention over chunks of 512 tokens, and global self-attention in the [CLS] token.

As a second strategy, we considered a Chunk Encoding (CE) approach, dividing the documents into C chunks and processing them individually with GatorTron-base. By dividing inputs into chunks, we can effectively leverage the capabilities of a standard Transformer encoder, limited to a maximum of T tokens (in our case, $T = 512$), to analyze long clinical documents. While both LE and CE are simple to implement, CE can perhaps better avoid difficulties in having LMs making robust use of information within long inputs (Liu et al., 2023).

CE relies on the assumption that if an ICD code can be identified in a single segment (i.e., a chunk) of the input document, then the code should naturally also be assigned when classifying the document as a whole (i.e., a single mention should be




Figure 1: A simple classification architecture that considers the Chunk Encoding (CE) approach.




Figure 2: Smooth document segmentation with 255 token overlaps. Each chunk includes, at the end, the sentence separation token [SEP] characteristic of BERT-type models, completing 512 tokens per chunk.

enough to justify the coding decision). We thus use a max-pooling operation to consolidate the detection of different ICD codes in each chunk, as illustrated in Figure 1, where C refers to the number of chunks, T corresponds to the number of tokens within each chunk, H corresponds to the dimensionality of the vectors representing each token, and L denotes the number of ICD classes. We also adopted a smooth partitioning scheme that considers overlaps between chunks to mitigate the loss of information from abruptly breaking interconnected pieces of text, as shown in Figure 2.

In the remaining parts of the paper, we will collectively refer to GatorTron and its Longformer version as the GatorTron encoder (GatorEnc). Note that when using the LE strategy, the number of chunks C can be seen as being equal to one, and we do not require the max-pooling operation to consolidate chunk results. With this, we can interpret the following images and expressions in a way that generalizes to both the LE and CE approaches.

3.2 Multi-Synonyms Attention

Inspired by Yuan et al. (2022), we enhanced our classification model by integrating a multi-synonyms attention mechanism. The primary objective was to explore the intricate relationships between specific mentions of ICD codes, within chunks of the hospital discharge summaries, and the textual descriptions of ICD codes. This integration aimed to leverage synonyms to improve code

³<https://catalog.ngc.nvidia.com/>

⁴<https://huggingface.co/UFNLP/gatortron-base>

representation learning (i.e., label embeddings), aiding in code classification.

First, we extended the ICD code descriptions with synonyms obtained from a large medical knowledge base, specifically the UMLS⁵ metathesaurus. By aligning ICD codes with UMLS Concept Unique Identifiers (CUIs), we selected corresponding synonyms for English terms sharing the same CUIs. Additionally, we considered synonym variants by removing special characters, allowing only hyphens and brackets, and removing the coordinating conjunctions *or* and *and*.

To improve diversity, we also gathered additional synonyms from Wikidata and Wikipedia. However, even with this addition, we observed that the lists of synonyms associated with each ICD code were often repetitive, posing a risk of introducing undesired bias in classification. We selected a maximum of M synonyms by first representing them as vectors through GatorEnc (i.e., taking the [CLS] token representation for each synonym). Then, a set of M diverse vectors were selected for each ICD code through the application of the Gurobi optimizer⁶ as a way to address the Maximum Diversity Problem⁷ (MDP), introduced by Glover et al. (1977) and which can be formulated as follows:

$$\text{maximize } \sum_{i=1}^{N-1} \sum_{j=i+1}^N d_{ij} x_i x_j, \quad (1)$$

$$\text{subject to } \sum_{i=1}^N x_i = M, \quad (2)$$

$$x_i \in \{0, 1\}, \quad 1 \leq i \leq N. \quad (3)$$

In the previous equations, d_{ij} is a distance metric between synonym representations i and j (i.e., the cosine distance between the vectors), and x_i takes the value one if element i is selected and 0 otherwise. By solving the MDP, we select a small subset of M out of N synonyms, that effectively represent the broader embedding space for each ICD code.

We can denote by Q_l a matrix where rows correspond to the representations for the M synonyms associated to an ICD code l , with each code synonym j_l composed of tokens $\{s_i^{j_l}\}_{i=1}^{S_{j_l}}$:

$$Q_l = \{\text{GatorEnc}(s_1^{j_l}, \dots, s_{S_{j_l}}^{j_l})[\text{CLS}]\}_{j_l=1}^M. \quad (4)$$

Note that the synonym representations are not updated during model training. The token representations of hidden size H , within each chunk of text c , are similarly produced with GatorEnc, as follows:

$$K^c = \text{GatorEnc}(x_1^c, \dots, x_T^c). \quad (5)$$

⁵<https://www.nlm.nih.gov/research/umls/>

⁶<https://www.gurobi.com>

⁷<https://grafo.etsii.urjc.es/optsi.com/mdp.html>

Thus, K corresponds to the aggregate token representations of all chunks. To integrate the text representations from each chunk with the multiple synonym representations, we use an approach inspired by the multi-synonyms attention method proposed by Yuan et al. (2022), which in turn draws inspiration from the multi-head attention mechanism of the Transformer architecture (Vaswani et al., 2017). We specifically split each K^c into Z heads, equaling this value to the maximum number of synonyms per code, i.e., $Z = M$:

$$K^c = K_1^c, \dots, K_Z^c. \quad (6)$$

The code synonyms $\{Q_l\}_{l=1}^L$ are used to query each K^c and, by calculating attention scores α_l over K^c , we identify the parts from the chunk’s text that are more related to code synonym l :

$$\alpha_l = \{\text{Softmax}(W_Q Q_l \cdot \text{Tanh}(W_K K^c))\}_{c=1}^C. \quad (7)$$

We use an average pooling operation over $\text{Tanh}(K)\alpha_l$ to create code-wise text representations R , averaging the contributions from synonyms:

$$R = \{\text{AvgPool}(\text{Tanh}(K)\alpha_l)\}_{l=1}^L. \quad (8)$$

To assess whether the text of a chunk c contains code l , we evaluate the similarity between the code-wise text representation R_c and the code’s embeddings V . We aggregate the synonym representations $\{Q_l\}_{l=1}^L$ to form code representations V through average pooling, resulting in a matrix with each row depicting a global code representation:

$$V = \{\text{AvgPool}(Q_l^1, Q_l^2, \dots, Q_l^M)\}_{l=1}^L. \quad (9)$$

To measure the similarity for classification, we apply a bi-affine transformation. Finally, after carefully attending to the ICD codes in each chunk, using synonyms to enhance the classification, we employ max-pooling to consolidate the results:

$$Y = \sigma(\text{MaxPool}(\text{Diag}(R_1^t W V), \dots, \text{Diag}(R_C^t W V))). \quad (10)$$

In the previous equation, R_c^t is the transpose of R_c . Unlike previous approaches that perform classification using learned code-dependent parameters, which can be challenging to define for rare codes, our bi-affine function uses the parameters WV , where W is learned. This simplifies the learning process, at the same time making it more effective.

Figure 3 illustrates the combination of the chunk-based encoding strategy, described in the previous

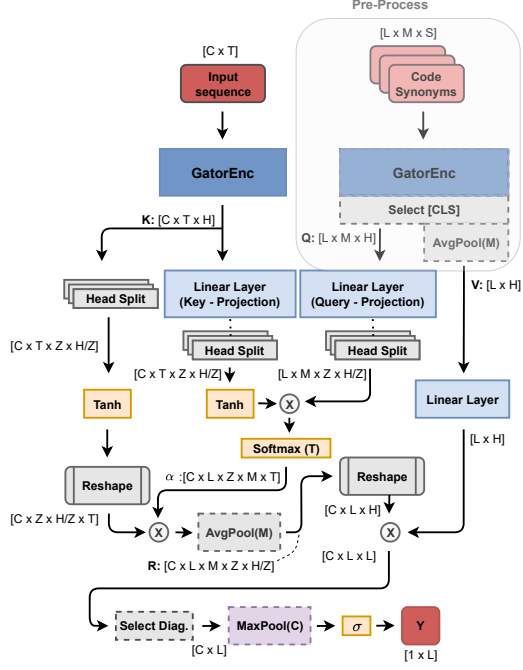


Figure 3: The classification architecture that combines the CE with a multi-synonyms attention mechanism.

section, with the classification method that considers the multi-synonyms attention mechanism.

For model training, noting that we are in the presence of a multi-label classification task, we adopted the widely-used Binary Cross-Entropy (BCE) loss, which treats each class independently and can be formally described as follows:

$$\mathcal{L}_C = \sum_{l=1}^L -y_l \log(\hat{y}_l) - (1 - y_l) \log(1 - \hat{y}_l). \quad (11)$$

The variable $y_l \in \{0, 1\}$ represents the ground-truth for a code l , while \hat{y}_l represents the probability of that code being present, as given by the classifier, and L is the number of different ICD codes.

3.3 Joint Classification and Quantification

Following previous work by Coutinho and Martins (2023), we considered an approach inspired by under-complete denoising auto-encoders to quantify the prevalence of ICD codes within a set of documents, accounting with label associations. We integrated this quantification module, implemented as a three-layer MLP, together with the classifier, performing end-to-end training of the resulting model. We hypothesize that the classification and the quantification objectives can naturally complement each other, and that combining them can contribute to a better calibration of the model.

Notice that classification operates at the level of individual instances, while quantification operates over groups of instances. To integrate both

objectives within end-to-end training, we follow the steps described next:

1. **Shuffling and setting a limit:** We shuffle the dataset at the start of each training epoch. We also establish a limit that simulates the maximum number of instances to be considered for quantification, selected randomly between one and the total number of training instances.
2. **Iterative data collection:** We process the instances individually as we progress through the training set. We collect the classification results for each instance that is processed, until we hit the previously defined maximum limit. This creates a new group of instances for each instance that is processed, consisting of all the instances that we have processed thus far, plus the latest one. The processing of each instance is made as follows:

(a) **Computation of classification loss:**

When processing each new instance, we apply our classification model and calculate the classification loss associated with that instance.

(b) **Computation of quantification loss:**

We add the classification output to the previous classification outputs for instances within the group. This allows us to compute a Probabilistic Classify and Count (PCC) vector, denoting the estimated relative frequency of each class label within the group of instances. We then process this vector using the aforementioned MLP, which refines the PCC estimates. We finally calculate the quantification loss with the refined estimates.

(c) **Aggregation of results:**

The loss values computed in the previous steps are aggregated into a total loss, which is used to update the model parameters for each batch of processed instances.

3. **Repeat and reset:** We follow the iterative process (i.e., steps (a) to (c)) until we reach the maximum number of instances designated for the quantification set. Once this limit is reached, we reset the quantification group and establish a new maximum limit for quantifying instances. We continue with model training until a stopping criterion is met.

The proposed approach performs joint training with classification and quantification, accumulating

the PCC quantification estimates over batches of instances to produce refined quantification results for groups of different sizes. At inference time, we can take the MLP that was trained jointly, and use it separately to perform quantification (i.e., a classifier can process each instance in a group in order to produce class probabilities, which can be aggregated into a PCC quantification estimate to be then refined with the MLP).

Our combined loss function can be formally described by the following equation, where λ is a hyper-parameter controlling the relative influence of the quantification loss:

$$\mathcal{L} = \mathcal{L}_C + \lambda \mathcal{L}_Q. \quad (12)$$

The classification loss (\mathcal{L}_C) is the BCE, formally described in Equation 11, and the quantification loss (\mathcal{L}_Q) can correspond to either the common Mean Squared Error (MSE) loss (\mathcal{L}_Q^{MSE}), or the Huber loss (\mathcal{L}_Q^{Huber}), respectively given by:

$$\mathcal{L}_Q^{MSE}(\hat{p}_\epsilon, p_\epsilon) = \sum_{l=1}^L (\Delta_l)^2, \quad (13)$$

$$\mathcal{L}_Q^{Huber}(\hat{p}_\epsilon, p_\epsilon) = \begin{cases} \frac{1}{2} \sum_{l=1}^L (\Delta_l)^2 & \text{if } \Delta_l < \delta, \\ \delta \cdot (\Delta_l - \frac{1}{2}\delta) & \text{otherwise.} \end{cases} \quad (14)$$

In the previous expressions, $\Delta_l = |\hat{p}_\epsilon(l) - p_\epsilon(l)|$, where p_ϵ refers to the ground-truth quantification result (i.e., the relative class frequency within the set of instances) for each of the L class labels, \hat{p}_ϵ refers to the quantification estimates, and δ is a tuning parameter that determines the point at which the Huber loss transitions from a quadratic to a linear penalty. The Huber loss has a smooth optimization landscape, like the MSE, but it is less sensitive to outliers, thus perhaps leading to more stable and accurate results.

4 Experimental Evaluation

This section presents the experimental evaluation of the proposed method, establishing a comparison with previously reported results.

4.1 Datasets

Experiments were conducted using the publicly available MIMIC-III data (Johnson et al., 2016), which we accessed through PhysioNet⁸ after completing the ethical training program from the associated collaborative institutional training initiative. We specifically used the same dataset splits

⁸<https://physionet.org/content/mimiciii/>

considered in previous studies, namely MIMIC-III-50 (Mullenbach et al., 2018), which only comprises the top-50 most frequent ICD-9 codes in the dataset, and also MIMIC-III-clean (Edin et al., 2023), which corresponds to a cleaned version of the dataset that contains 3,681 unique ICD-9 codes. We present a statistical characterization of the dataset splits in Appendix A.1.

The quantification experiments also used MIMIC-III-50 and MIMIC-III-clean, following the general methodology from Coutinho and Martins (2023). Specifically, to assess result quality in each case, we sampled documents from the corresponding validation set to form 5,000 quantification groups of different sizes, with the size parameter varying between one and the number of documents in the set. A separate set of 1,000 groups was also created by sampling documents from the test split. These sets were used for model training and testing in the quantification experiments.

4.2 Evaluation Metrics

We assessed the proposed approach across various metrics considered in previous work, to ensure a fair comparison with prior research.

Regarding the classification task, we used micro and macro-averaged F1 scores, Area Under the Curve (AUC) scores, and precision at cutoff n . For the experiments over the MIMIC-III-50 dataset, we defined $n = 5$. For the experiments conducted on MIMIC-III-clean, we considered $n = 8$ and $n = 15$, roughly aligning with the average number of codes in each split. To measure our classifier’s calibration quality, we used the Mean Expected Calibration Error (MECE) with 20 bins.

To evaluate the quantification task, we used the Mean Absolute Error (MAE) and the Mean Relative Absolute Error (MRAE) (Coutinho and Martins, 2023). The MRAE uses an additive smoothing penalty to avoid divisions by zero, which slightly impacts results according to the size of the groups.

4.3 Implementation Details

Table 1 presents the training hyper-parameters considered in our experiments.

Models using the Chunk Encoder (CE) have a maximum allowed number of input tokens limited only by hardware constraints. During training, we had to limit this input length according to the available GPU memory, considering a single NVIDIA A100 with 80Gb. However, at inference time, we could increase this limit up to 20,000 tokens. In

Parameters	MIMIC-III-50	MIMIC-III-clean
Maximum token input length for CE	7,142	6,122
Maximum token input length for LE	8,192	8,192
Token overlapping window for CE	255	255
GatorEnc hidden size (H)	1,024	1,024
Synonyms per ICD code (M)	4	4
Number of heads (Z)	4	4
Maximum number of epochs	300	300
Early stopping patience	5	5
Effective batch size	16	16
Adam ϵ	1e-08	1e-08
Starting learning rate	2e-05/2e-07	2e-05/2e-07
Ending learning rate	0	0
MLP hidden size	32	3,072
Quantification coefficient (λ)	100	100

Table 1: Hyper-parameters used for model training in the MIMIC-III-50 and MIMIC-III-clean settings. The *maximum number of epochs* values are related to both the classification and quantification modules.

turn, models using the Longformer Encoder (LE) have the same maximum input length for training and testing. Once again, we set this value based on the available GPU memory.

We trained our classifiers in two stages, using a linear scheduler for the learning rate. The first stage uses a learning rate starting at $2e-5$ and proceeds until we reach an early stopping criteria based on the micro-averaged F1 score over the development set. We then perform a second training stage, with a learning rate starting at $2e-7$ and an early stopping criteria based on the Mean Expected Calibration Error (MECE). The quantification model (i.e., the MLP) was first trained individually following the guidelines of [Coutinho and Martins \(2023\)](#), using a learning rate that starts at $2e-5$ and proceeding until we reach an early stopping criterion, based on the MSE loss. The model that integrates the quantification objective was initialized with pre-trained classification and quantification components obtained after a first stage of training. Thus, these components should already perform each task with reasonable competence, prior to their combination.

4.4 Experiments and Results

We comprehensively evaluated the proposed approach with different metrics, comparing it against previous methods and ablated model versions.

Our Baseline Model (BM) uses GatorTron-base to process the first 512 tokens of each document, without the Longformer Encoder (LE) or Chunk Encoder (CE) strategies, and without the Multiple-Synonyms Attention Mechanism (MSAM). We also assessed the combination of the two encod-

ing strategies with the label embeddings (i.e., models referred to as {CE, LE}+MSAM), and also the joint training with classification and quantification objectives (i.e., {CE, LE}+MSAM+CLQ).

4.4.1 Classification

Tables 2 and 3 present classification results for different model variants, respectively for MIMIC-III-50 and MIMIC-III-clean.

Both encoding strategies (CE and LE) significantly outperform the Baseline Model (BM), with CE also outperforming LE in all metrics by a considerable margin. Notice that LE compresses the entire document into a single representation vector, while the CE strategy considers identifying ICD codes in smaller chunks, avoiding problems in encoding long inputs ([Liu et al., 2023](#)).

The MSAM mechanism also notably enhances performance across all the metrics. Still, despite the significant performance gap between the LE and CE models, these differences diminish after incorporating this module. In MIMIC-III-clean, LE+MSAM has a slightly higher micro-averaged F1 score (+0.6) but a significantly lower macro-averaged F1 score (-2.5) when compared to CE+MSAM. This indicates that while both versions show similar overall results across the different ICD classes, CE performs better in classifying individual ICD classes with lower prevalence frequencies, which can be important in this domain.

Considering the overall better results with the CE strategy, we decided to use only this encoding approach for further experiments, considering joint training with the quantification objective. The results in Tables 2 and 3 show that the joint training, either with the MSE or the Huber loss, does not significantly impact the classification accuracy.

To assess the impact of using a different number of synonyms in the label embeddings, and also the diversity-based strategy for selecting the synonyms, we considered the CE+MSAM model over the MIMIC-III-50 dataset. We varied M between 2, 4, or 8 synonyms, and considered either random or maximum-diversity (i.e., MDP) selection. The results are shown in Table 4, confirming that the maximum-diversity selection strategy positively impacts the results when using fewer synonyms. Consistently with the results from [Yuan et al. \(2022\)](#), our experiments also indicate that $M = 4$ produces the best results.

We also analyzed the proposed approach in terms of calibration performance. In Table 5, we explic-

Model	Stopping Epochs	AUC		F1		P@N	
		Macro	Micro	Macro	Micro	P@5	P@5
CAML* (Mullenbach et al., 2018)	–	87.5	91.1	51.0	60.6	61.1	
MSATT-KG [†] (Xie et al., 2019)	–	91.4	93.6	63.8	68.4	64.4	
MultiResCNN* (Li and Yu, 2020)	–	89.7	92.4	61.1	67.3	64.4	
LAAT* (Vu et al., 2020)	–	90.5	92.8	59.2	66.8	64.0	
PLM-ICD* (Huang et al., 2022)	–	91.7	93.8	65.4	70.5	65.7	
MSMN [†] (Yuan et al., 2022)	–	92.8	94.7	68.3	72.5	68.0	
KEPTLongformer [†] (Yang et al., 2022b)	–	92.6	94.8	68.9	72.9	67.3	
BM	11(+0)	83.8	87.0	49.1	56.1	55.8	
LE	26(+0)	84.9	87.9	54.8	61.5	59.7	
CE	10(+0)	91.2	93.4	65.5	70.0	66.1	
LE+MSAM	5(+6)	93.8	95.4	70.3	73.9	69.1	
CE+MSAM	4(+10)	93.7	95.4	70.4	73.9	68.8	
CE+MSAM+CLQ _{MSE}	4(+4)	93.7	95.4	70.4	74.0	68.9	
CE+MSAM+CLQ _{Huber}	4(+1)	93.7	95.4	70.3	73.8	68.9	

Table 2: Results for the different classification methods on the MIMIC-III-50 test set. Results for methods marked with * were taken directly from Edin et al. (2023). Results for methods marked with † were taken directly from the corresponding paper. The values in bold represent the best-in-class performance in terms of the different evaluation metrics.

Model	Stopping Epochs	AUC		F1		P@N	
		Macro	Micro	Macro	Micro	P@8	P@15
CAML* Mullenbach et al. (2018)	–	91.4	98.2	20.4	55.4	67.7	52.8
MultiResCNN* (Li and Yu, 2020)	–	93.1	98.5	22.9	56.4	68.5	53.5
LAAT* (Vu et al., 2020)	–	94.0	98.6	22.6	57.8	70.1	54.8
PLM-ICD* (Huang et al., 2022)	–	95.9	98.9	26.6	59.6	72.1	56.5
BM	43(+0)	89.9	95.7	11.0	44.5	59.5	44.2
LE	79(+0)	90.3	95.7	12.9	48.6	63.0	46.9
CE	68(+0)	91.7	96.1	16.9	52.1	66.1	50.6
LE+MSAM	6(+6)	96.3	99.0	28.0	60.9	74.2	58.1
CE+MSAM	8(+5)	96.3	98.9	30.5	60.3	73.3	57.5
CE+MSAM+CLQ _{MSE}	8(+6)	96.4	99.0	31.2	60.5	73.3	57.4
CE+MSAM+CLQ _{Huber}	8(+5)	96.3	98.9	30.5	60.4	73.3	57.4

Table 3: Results for the different classification methods on the MIMIC-III-clean test set. Results for methods marked with * were taken from Edin et al. (2023).

itly examine the calibration error over different sets of ICD codes: Low percentile (*Low Pth*) corresponds to the average value of the calibration error calculated for the 10% of ICD codes with the lowest frequency rates in the training set of the respective MIMIC-III split. In turn, the medium percentile (*Medium Pth*) represents the average value of the calibration error for the 10% of ICD codes with medium frequency rates, falling within the 55% to 65% percentile range in the respective MIMIC-III split training set. Finally, the high percentile (*High Pth*) indicates the average value of the calibration error for the 10% of medical codes with the highest frequency of occurrence in the training set of the respective MIMIC-III split.

The results show that the label embedding mechanism offers notable benefits in model calibration. The joint optimization also improved calibration results for MIMIC-III-50, although not in MIMIC-III-clean, where LE+MSAM exhibited the best average calibration performance. Chunk-based mod-

	AUC		F1		P@N
	Macro	Micro	Macro	Micro	P@5
M = 1	93.5	95.2	69.3	72.5	68.0
M = 2 (random)	93.3	95.2	69.4	72.8	68.2
M = 2 (maximum-diversity)	93.6	95.3	69.8	73.4	68.3
M = 4 (random)	93.6	95.3	69.8	73.3	68.4
M = 4 (maximum-diversity)	93.7	95.4	70.4	73.9	68.8
M = 8 (random)	93.4	95.1	69.9	73.4	68.2
M = 8 (maximum-diversity)	93.4	95.1	69.2	72.9	68.0

Table 4: Results when considering a different number of synonyms (M) on the MIMIC-III-50 dataset.

Dataset	Classifier	Mean	Low Pth	Medium Pth	High Pth
MIMIC-III-50	BM	3.8e-02	1.7e-02	3.0e-02	6.1e-02
	LE	6.0e-02	3.5e-02	6.4e-02	8.2e-02
	CE	3.5e-02	2.1e-02	3.0e-02	5.1e-02
	LE+MSAM	2.7e-02	2.0e-02	2.5e-02	3.6e-02
	CE+MSAM	2.8e-02	1.9e-02	2.4e-02	3.7e-02
	CE+MSAM+CLQ _{MSE}	2.9e-02	2.1e-02	2.6e-02	3.7e-02
CE+MSAM+CLQ _{Huber}	2.6e-02	1.9e-02	2.5e-02	3.3e-02	
MIMIC-III-clean	BM	218.1e-05	8.8e-05	60.0e-05	1520.0e-05
	LE	258.2e-05	12.2e-05	90.0e-05	1660.8e-05
	CE	248.6e-05	11.7e-05	87.8e-05	1595.1e-05
	LE+MSAM	140.1e-05	18.4e-05	81.2e-05	680.1e-05
	CE+MSAM	155.6e-05	18.5e-05	85.3e-05	769.0e-05
	CE+MSAM+CLQ _{MSE}	161.4e-05	20.1e-05	87.7e-05	800.6e-05
CE+MSAM+CLQ _{Huber}	157.2e-05	18.3e-05	85.3e-05	780.6e-05	

Table 5: Calibration quality according to the MECE metric, for all the proposed classification models and on different percentiles of the MIMIC-III splits.

eling can perhaps overvalue the probabilities associated to each class, negatively affecting the calibration with its use of a max-pooling operation. Additionally, our results also show that the Huber loss has a positive effect on enhancing model calibration, compared to the MSE loss. Although the jointly trained model that optimizes quantification through the MSE loss has slightly better classification results, the variant that uses the Huber loss is better in terms of calibration. We, therefore, used the CE+MSAM+CLQ_{Huber} model for further considerations, arguing that this variant can better balance classification and calibration performance.

Compared to other approaches in the literature, we outperform the previously best-performing models reported for the two MIMIC-III splits under analysis. It is worth noting that the models reported by Edin et al. (2023) underwent an adjustment using the validation splits, as the authors reported on classification performance after optimizing the decision boundary values through a grid search mechanism to maximize F1 scores in the validation splits. In contrast, our results do not involve any such adjustment and still surpass the best-reported models to date, establishing a new state-of-the-art with a default decision boundary set at 0.5.

On MIMIC-III-50, the proposed approach outperforms the best reported model to date (i.e.,

KEPTLongFormer) across all metrics, with scores of 93.7 (+1.1), 95.4 (+0.6), 70.3 (+1.4), 73.8 (+0.9), and 68.9 (+1.6) in terms of macro-AUC, micro-AUC, macro-F1, micro-F1, and P@5, respectively. On MIMIC-III-clean, we outperform the best reported model to date (i.e., PLM-ICD) also across all metrics, with scores of 96.3 (+0.4), 98.9 (+0.0), 30.5 (+3.9), 60.4 (+0.8), 73.3 (+1.2) and 57.4 (+0.9) in terms of macro-AUC, micro-AUC, macro-F1, micro-F1, P@8, and P@15.

Appendix A.2 details the classification performance across different ICD chapters, additionally also showing results for the top-10 most frequent ICD codes, and for relevant chronic diseases. These different examples attest to the usefulness of our approach, offering accurate classification results that can inform different types of downstream analyses.

4.4.2 Quantification

Tables 6 and 7 show results for quantification experiments using both MIMIC-III splits. The baseline results correspond to the standard CC and PCC methods, and also to the use of an MLP separately trained for quantification, following the experimental setup from Coutinho and Martins (2023). In the case of the proposed models, i.e. CE+MSAM+CLQ_{MSE} and CE+MSAM+CLQ_{Huber}, the MLP trained jointly with the classifier, using either the MSE or the Huber loss, was employed for quantification.

CE+MSAM+CLQ_{MSE} outperforms CE+MSAM and CE+MSAM+CLQ_{Huber} in terms of classification accuracy, but not in terms of calibration, which impacts the standard CC and PCC metrics. Notably, the LE+MSAM model achieves better results with PCC but worse with CC, aligning with the idea that chunk encoding tends to overvalue class probabilities, negatively impacting the PCC results. Still, this effect is mitigated when not looking at the probabilities directly, i.e. with the CC method.

Analyzing Tables 6 and 7, we see that joint optimization generally surpasses all mentioned baselines, including the separate training of the MLP proposed by Coutinho and Martins (2023). This is consistent across all metrics, except for the MRAE in MIMIC-III-clean, where combining BM with PCC corresponds to the lowest value. Although this result seems misleading, the MRAE uses additive smoothing to avoid divisions by zero, which slightly impacts results according to group sizes. Is it thus important to also look at results with the MAE, i.e., a complementary metric that evaluates

Model	CC		PCC		MLP/CLQ	
	MAE	MRAE	MAE	MRAE	MAE	MRAE
BM	4.08e-02	22.02e-02	1.70e-02	10.12e-02	1.22e-02	7.22e-02
LE	2.83e-02	15.70e-02	2.02e-02	11.53e-02	1.15e-02	6.87e-02
CE	2.11e-02	10.08e-02	1.50e-02	9.67e-02	1.14e-02	6.83e-02
LE+MSAM	1.96e-02	10.11e-02	1.25e-02	7.34e-02	1.09e-02	6.64e-02
CE+MSAM	1.72e-02	9.31e-02	1.38e-02	9.31e-02	1.09e-02	6.64e-02
CE+MSAM+CLQ _{MSE}	1.91e-02	9.90e-02	1.69e-02	10.90e-02	1.14e-02	6.83e-02
CE+MSAM+CLQ _{Huber}	1.89e-02	10.32e-02	1.47e-02	10.0e-02	1.05e-02	6.64e-02

Table 6: Results for different quantification methods, using the results from different classification models on the MIMIC-III-50 test dataset split.

Model	CC		PCC		MLP/CLQ	
	MAE	MRAE	MAE	MRAE	MAE	MRAE
BM	18.73e-04	3.03e-01	8.39e-04	1.88e-01	8.07e-04	2.18e-01
LE	16.00e-04	2.86e-01	9.99e-04	2.01e-01	8.00e-04	2.18e-01
CE	15.78e-04	2.76e-01	11.23e-04	2.02e-01	7.89e-04	2.12e-01
LE+MSAM	12.14e-04	2.84e-01	7.34e-04	2.21e-01	7.83e-04	2.13e-01
CE+MSAM	11.48e-04	2.27e-01	7.78e-04	2.21e-01	7.68e-04	2.09e-01
CE+MSAM+CLQ _{MSE}	11.13e-04	2.26e-01	9.12e-04	2.55e-01	7.00e-04	2.04e-01
CE+MSAM+CLQ _{Huber}	11.41e-04	2.28e-01	8.12e-04	2.30e-01	6.93e-04	2.04e-01

Table 7: Results for different quantification methods, using the results from different classification models on the MIMIC-III-clean test dataset split.

quantification consistently across group sizes.

Overall, the results indicate that joint training, and particularly when using the Huber loss, can effectively inform the MLP about class distributions, leading to good quantification performance and improving classifier calibration compared to the MSE loss. We present a more detailed analysis of the quantification results in Appendix A.3.

5 Conclusions and Future Work

This work introduced a novel deep learning method for ICD coding, achieving state-of-the-art results in tests with two well established MIMIC-III dataset splits. The proposed method processes long clinical documents and uses a label embedding mechanism that explores diverse ICD code synonyms. Besides achieving highly accurate classification results, the proposed approach produces well-calibrated estimates that can effectively inform downstream tasks such as text quantification.

Despite the strong results, it should be noted that our model does not exploit the hierarchical structure inherent to the ICD coding system. Thus, a promising avenue for further improvement involves using this structural knowledge, e.g. by implementing dual classification heads. Another path worth exploring relates to using alternative methods to improve calibration (e.g., using other classification loss functions besides the BCE), since improving calibration is beneficial for classification and essential for accurate results in quantification.

Acknowledgements

We thank the anonymous reviewers for their valuable comments and suggestions. This research was supported by the Portuguese Recovery and Resilience Plan through project C645008882-00000055 (i.e., the Center For Responsible AI), and also by Fundação para a Ciência e Tecnologia (FCT), through the projects with references DSAIPA/DS/0133/2020 and UIDB/50021/2020 (DOI:10.54499/UIDB/50021/2020).

Limitations and Ethical Considerations

While our work does not raise new ethical issues within this domain, there are general concerns.

ICD coding is essential in clinical, operational, and financial healthcare decisions. Traditionally, medical coders review documents and manually assign the appropriate ICD codes by following specific coding guidelines. Approaches such as ours can significantly reduce time and costs in ICD coding. Still, there are risks associated with over-reliance on automatic coding methods. No matter how accurate a given approach is, it is still possible to misclassify documents with erroneous ICD codes, affecting patient treatment. Therefore, automatic coding should assist, rather than replace, the judgment of trained clinical professionals.

Our experiments have also relied on MIMIC-III dataset splits used in previous studies. While these datasets constitute useful benchmarks for developing and evaluating new methods, they do not represent the enormous variety of clinical and linguistic data encountered in potential deployments.

References

- Iz Beltagy, Matthew E. Peters, and Arman Cohan. 2020. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*.
- Isabel Coutinho and Bruno Martins. 2023. Exploring label correlations for quantification of ICD codes. In *Proceedings of the International Conference on Discovery Science*.
- Xiang Dai, Ilias Chalkidis, Sune Darkner, and Desmond Elliott. 2022. Revisiting transformer-based models for long document classification. *arXiv preprint arXiv:2204.06683*.
- Joakim Edin, Alexander Junge, Jakob D. Havtorn, Lasse Borgholt, Maria Maistro, Tuukka Ruotsalo, and Lars Maaløe. 2023. Automated medical coding on MIMIC-III and MIMIC-IV: A critical review and replicability study. In *Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval*.
- Fred Glover, G. Hersh, and C. McMillian. 1977. Selecting subsets of maximum diversity. Technical Report MS/IS 77-9, University of Colorado at Boulder.
- Chao-Wei Huang, Shang-Chi Tsai, and Yun-Nung Chen. 2022. PLM-ICD: Automatic ICD coding with pretrained language models. *arXiv preprint arXiv:2207.05289*.
- Alistair E. W. Johnson, Tom J. Pollard, Lu Shen, Liwei 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*, 3:1–9.
- Joon Lee, Daniel J. Scott, Mauricio Villarroel, Gari D. Clifford, Mohammed Saeed, and Roger G. Mark. 2011. Open-access MIMIC-II database for intensive care research. In *Proceedings of the International Conference of the IEEE Engineering in Medicine and Biology Society*.
- Fei Li and Hong Yu. 2020. ICD coding from clinical text using multi-filter residual convolutional neural network. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2023. Lost in the middle: How language models use long contexts. *arXiv preprint arXiv:2307.03172*.
- Alejandro Moreo, Manuel Francisco, and Fabrizio Sebastiani. 2022. Multi-label quantification. *arXiv preprint arXiv:2211.08063*.
- James Mullenbach, Sarah Wiegrefe, Jon Duke, Jimeng Sun, and Jacob Eisenstein. 2018. Explainable prediction of medical codes from clinical text. *arXiv preprint arXiv:1802.05695*.
- Kimberly J. O'Malley, Karon F. Cook, Matt D. Price, Kimberly Raiford Wildes, John F. Hurdle, and Carol M. Ashton. 2005. Measuring diagnoses: ICD code accuracy. *Health Services Research*, 40:1620–1639.
- Tobias Schumacher, Markus Strohmaier, and Florian Lemmerich. 2021. A comparative evaluation of quantification methods. *arXiv preprint arXiv:2103.03223*.
- 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 *Proceedings of the Annual Conference on Advances in Neural Information Processing Systems*.
- Thanh Vu, Dat Quoc Nguyen, and Anthony Nguyen. 2020. A label attention model for ICD coding from clinical text. *arXiv preprint arXiv:2007.06351*.

- Xiancheng Xie, Yun Xiong, Philip S. Yu, and Yangyong Zhu. 2019. EHR coding with multi-scale feature attention and structured knowledge graph propagation. In *Proceedings of the ACM International Conference on Information and Knowledge Management*.
- Xi Yang, Aokun Chen, Nima PourNejatian, Hoo Chang Shin, Kaleb E. Smith, Christopher Parisien, Colin Compas, Cheryl Martin, Mona G. Flores, Ying Zhang, et al. 2022a. GatorTron: A large clinical language model to unlock patient information from unstructured electronic health records. *arXiv preprint arXiv:2203.03540*.
- Zhichao Yang, Shufan Wang, Bhanu Pratap Singh Rawat, Avijit Mitra, and Hong Yu. 2022b. Knowledge injected prompt based fine-tuning for multi-label few-shot ICD coding. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*.
- Zheng Yuan, Chuanqi Tan, and Songfang Huang. 2022. Code synonyms do matter: Multiple synonyms matching network for automatic ICD coding. *arXiv preprint arXiv:2203.01515*.

A Appendix

This appendix presents statistical information about the dataset splits, and additional experimental results for the classification and quantification tasks.

A.1 MIMIC-III Dataset Splits

Table 8 provides a statistical characterization of the MIMIC-III splits considered in our experiments, detailing the training, validation, and test sets, and underlining the highly imbalanced label distribution, the disparity between the average and maximum document lengths, and the high number of ICD codes assigned to each discharge summary.

In turn, Table 9 presents the frequency of ICD codes over the data, dividing the ICD codes into three relevant percentiles for the training and test sets, for both MIMIC-III splits. Low Pth accounts for the 10% of medical codes with the lowest occurrence frequency in the training set of the respective MIMIC-III split. Medium Pth corresponds to the 10% of codes with medium occurrence frequency, falling within the 55% to 65% percentile range in the training set of the respective MIMIC-III split. Lastly, High Pth corresponds to the 10% of codes with the highest occurrence frequency in the training set of the respective MIMIC-III split.

Set (Split)	Samples	Words per Doc.		Tokens per Doc.		Codes per Doc.		Unique Codes	Type of Codes	
		Avg.	Max.	Avg.	Max.	Avg.	Max.		Diag.	Proc.
Train (top-50)	8,066	1,642	7,989	2,830	20,297	5.4	18	50	33	17
Val. (top-50)	1,573	1,932	6,658	3,410	16,566	5.9	21	50	33	17
Test (top-50)	1,729	1,964	6,470	3,465	11,871	6.0	20	50	33	17
Train (clean)	38,401	1,514	10,500	1,651	11,758	14.0	57	3,681	2,849	832
Val. (clean)	5,577	1,552	6,393	1,694	6,897	15.9	60	3,676	2,844	832
Test (clean)	8,734	1,485	7,858	1,619	8,299	14.8	56	3,681	2,849	832

Table 8: Statistics for the training, validation and test sets of MIMIC-III-50 (top) and MIMIC-III-clean (bottom). The columns labeled with *words per doc* refer to the average and maximum number of words per hospital discharge summary. *Tokens per doc* corresponds to the average and maximum number of tokens per clinical document. *Codes per doc* refers to the average and maximum number of ICD codes per document. *Unique codes* corresponds to the number of distinct ICD codes. Finally, *type of codes* is used to indicate the number of distinct diagnosis and procedure codes.

Dataset	Split	Low Pth	Medium Pth	High Pth
MIMIC-III-50	Train	397-449	759-914	1615-3233
	Test	60-127	148-247	402-470
MIMIC-III-clean	Train	4-9	36-56	308-14,598
	Test	1-4	6-27	55-2228

Table 9: Number of ICD code occurrences in specific percentiles of code prevalence frequency.

Code	Description	Precision	Recall	F1
401.9	Unspecified essential hypertension	76.73	85.20	80.75
38.93	Venous Catheterization, Not Elsewhere Classified	69.30	71.72	70.49
428.0	Heart failure	80.94	82.66	81.79
427.31	Atrial fibrillation	90.57	91.82	91.19
414.01	Coronary atherosclerosis of native coronary artery	81.17	86.35	83.68
96.04	Insertion Of Endotracheal Tube	79.78	80.99	80.38
96.6	Enteral Infusion Of Concentrated Nutritional Substances	70.24	76.86	73.40
584.9	Acute kidney failure, unspecified	73.70	68.79	71.16
250.00	Diabetes mellitus without mention of complication type II or unspecified type, not stated as uncontrolled	72.43	83.37	77.52
272.4	Other and unspecified hyperlipidemia	69.73	80.24	74.62
Average		76.46	80.8	78.50

Table 10: Results over the test split for the 10 most frequent ICD codes in the MIMIC-III-clean dataset.

Block	Chronic Disease	Unique codes (Present)	Percentage	Performance Metrics	
				Macro-F1	Micro-F1
250	Diabetes mellitus	33	1.943%	28.60	65.22
401-405	Hypertensive Disease	14	3.303%	29.10	76.84
410-414	Ischemic Heart Disease	32	3.279%	29.49	69.01
428	Heart Failure	15	2.471%	36.65	71.47
585;403-404	Renal Failure	16	1.600%	34.95	58.71
490-496	Pulmonary Disease	16	1.209%	45.05	67.56

Table 11: Results for relevant chronic diseases. The columns *unique codes* and *percentage* refer to the number of unique codes of the respective block within the MIMIC-III-clean dataset, and to the corresponding percentage of occurrences over the dataset.

A.2 Additional Classification Results

Tables 10 to 13 provide additional insights into our classification results, specifically considering the CE+MSAM+CLQ_{Huber} model.

Table 10 presents classification results for the top-10 most frequent ICD codes over MIMIC-III-clean. We obtained a mean precision of 76.46%, a mean recall of 80.8%, and a mean F1 score of 78.50%. In turn, Table 11 presents the performance for examples of relevant chronic diseases, representing some of the main focuses of healthcare investigation. Tables 12 and 13 provide results for codes within different ICD diagnosis and procedure chapters. Together, these results illustrate different possible applications for the ICD coding results.

Note that Chapter VII (i.e., *diseases of the circulatory system*) in the ICD-9 diagnosis codes accounts for a substantial portion of the MIMIC-III-clean dataset, representing 22.471% of all diagnosis codes. This chapter demonstrates good classification performance, with our model achieving a macro-averaged F1 score of 29.08% and a micro-averaged F1 score of 67.38%.

Conversely, Chapter XI (i.e., *complications of pregnancy, childbirth, and the puerperium*) is the

Chapter	Occurrences				Performance Metrics	
	Train	Validation	Test	Percentage	Macro-F1	Macro-F1
I	14,050	2,090	3,212	3.190%	33.57	52.91
II	9,200	1,401	2,076	2.09%	36.35	57.62
III	49,135	7,356	11,008	11.126%	32.99	60.33
IV	17,882	2,657	4,106	4.062%	29.94	40.93
V	17,392	2,562	3,740	3.905%	21.87	47.90
VI	15,811	2,433	3,397	3.567%	28.78	54.69
VII	99,076	14,729	22,526	22.471%	29.08	67.38
VIII	31,613	4,703	7,113	7.158%	35.46	95.54
IX	27,061	3,967	6,022	6.107%	30.59	56.32
X	22,940	3,438	5,260	5.215%	28.90	62.00
XI	151	24	33	0.034%	25.95	40.00
XII	6,056	888	1,371	1.371%	28.64	47.17
XIII	9,098	1,360	1,944	2.044%	27.66	51.08
XIV	2,228	328	471	0.499%	51.85	62.61
XV	12,656	1,740	2,565	2.796%	31.18	59.92
XVI	20,692	3,154	4,550	4.68%	15.32	39.50
XVII	87,280	13,018	19,131	19.685%	23.56	50.83

Table 12: Number of instances and classification performance metrics for each of the ICD diagnosis chapters. The column named *percentage* corresponds to the percentage of the diagnosis codes under consideration over the MIMIC-III-clean dataset.

Chapter	Occurrences				Performance Metrics	
	Train	Validation	Test	Percentage	Macro-F1	Macro-F1
I	5,508	855	1,347	3.589%	36.48	65.02
II	4,852	733	1,148	3.134%	41.18	66.73
III	91	13	17	0.056%	63.07	66.67
IV	102	15	23	0.065%	56.89	59.57
V	0	0	0	0%	0.0	0.0
VI	21	3	4	0.013%	40.00	40.00
VII	501	75	104	0.317%	24.58	36.59
VIII	9,590	1,480	2,164	6.161%	37.12	63.90
IX	47,762	6,895	10,813	30.478%	46.07	76.23
X	897	127	217	0.578%	49.38	71.36
XI	15,302	2,267	3,555	9.834%	38.18	66.54
XII	1,045	152	230	0.664%	54.41	74.89
XIII	641	102	127	0.405%	74.62	69.96
XIV	201	27	43	0.126%	56.65	65.06
XV	20	3	4	0.013%	88.89	88.89
XVI	5,990	924	1,307	3.827%	42.50	59.86
XVII	2,308	318	539	1.473%	32.19	50.41
XVIII	61,329	8,568	14,455	39.267%	25.37	66.70

Table 13: Number of instances and classification performance metrics for each of the ICD procedure chapters. The column named *percentage* corresponds to the percentage of the procedure codes under consideration over the MIMIC-III-clean dataset.

least frequent chapter of ICD codes, and corresponds to the lowest classification performance. With a prevalence of only 0.034% in the dataset, our model achieved macro- and micro-averaged F1 scores of 25.95% and 40.00%, respectively in this chapter. These scores highlight the negative impact of infrequent ICD codes on model effectiveness.

Furthermore, we observe an interesting phenomenon in Chapter XIV (i.e., *congenital anomalies*). Although this chapter represents a relatively small percentage (0.499%) of the overall dataset, the model performs remarkably well in this chapter. It attains macro and micro-averaged F1 scores of 51.85% and 62.61%, respectively, empirically showing the model’s ability to perform few-shot learning when dealing with seldom-seen codes.

When we examine the overall distribution of procedure codes, we see that the dataset is characterized by a generally low density of procedure codes, with two notable exceptions in Chapter IX (i.e., *operations on the cardiovascular system*) and Chapter XVIII (i.e., *miscellaneous diagnostic and therapeutic procedures*), which encompass almost 70% of the dataset. However, despite the relatively low frequency of procedures in the other chapters, our model performs exceptionally well in them. For instance, Chapters VI and XV achieve performance values of 40% and 88.89% respectively in macro- and micro-averaged F1, even though these codes have a small representation of 0.013% within the dataset. These results highlight the model’s capacity to learn even from infrequent instances, again emphasizing its few-shot learning capabilities.

Chapter XVIII in the ICD-9 procedure codes, which covers *miscellaneous diagnostic and therapeutic procedures*, stands out as the most frequently occurring chapter in the MIMIC-III-clean dataset, accounting for a substantial 39.267% of the total number of instances. In this chapter, we achieve 25.37% for macro-averaged F1 and 66.70% for micro-averaged F1.

A.3 Additional Quantification Results

Figure 4 shows Absolute Error (AE) and F1 scores per ICD class, sorted by prevalence frequency over the MIMIC-III-50 dataset.

The results show that the CLQ_{Huber} method outperforms PCC for nearly all ICD codes when it comes to accurately grasping ICD code prevalence. For instance, ICD code 401.9, which is the most frequent in the dataset, presents a high disparity in AE between PCC and CLQ_{Huber} results. An even higher disparity is seen for code 99.04, which corresponds to *transfusion of packed cells* and has an average prevalence frequency. Further investigation revealed that despite having a high F1 score (80.75%), the classification results for ICD code 401.9 have a notable difference between precision (76.73%) and recall scores (85.20%). This sug-

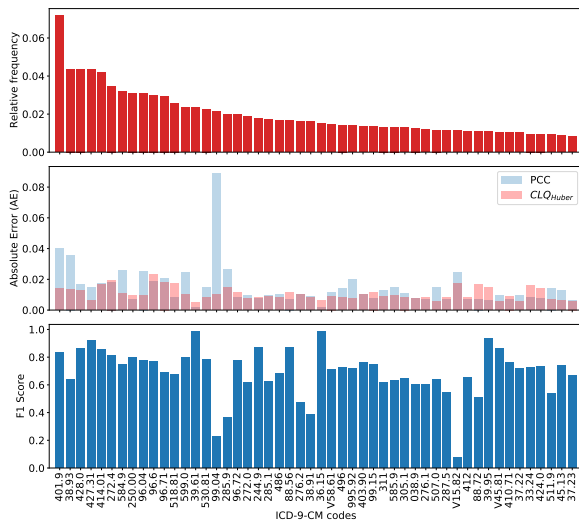


Figure 4: Relative frequency, absolute error, and F1 scores for each ICD code over MIMIC-III-50 dataset.




Figure 5: Estimated versus real prevalence for the two most frequent (top) and rarest (bottom) ICD codes in the MIMIC-III-50 dataset.

gests that the classification model overestimates this class, perhaps due to its high frequency, resulting in inaccurate posterior probabilities. The CLQ_{Huber} method appears to recognize this behavior and correct the results. Figure 5 aligns with the previous analysis, showing estimated versus real prevalence frequencies for frequent versus rare ICD codes (including code 401.9).