

Knowledge Injected Prompt Based Fine-tuning for Multi-label Few-shot ICD Coding

Zhichao Yang¹, Shufan Wang¹, Bhanu Pratap Singh Rawat¹, Avijit Mitra¹, Hong Yu^{1,2}

¹ College of Information and Computer Sciences, University of Massachusetts Amherst

² Department of Computer Science, University of Massachusetts Lowell

{zhichaoyang, shufanwang, brawat, avijitmitra}@umass.edu hong_yu@uml.edu

Abstract

Automatic International Classification of Diseases (ICD) coding aims to assign multiple ICD codes to a medical note with average length of 3,000+ tokens. This task is challenging due to a high-dimensional space of multi-label assignment (tens of thousands of ICD codes) and the long-tail challenge: only a few codes (common diseases) are frequently assigned while most codes (rare diseases) are infrequently assigned. This study addresses the long-tail challenge by adapting a prompt-based fine-tuning technique with label semantics, which has been shown to be effective under few-shot setting. To further enhance the performance in medical domain, we propose a knowledge-enhanced longformer by injecting three domain-specific knowledge: hierarchy, synonym, and abbreviation with additional pretraining using contrastive learning. Experiments on MIMIC-III-full, a benchmark dataset of code assignment, show that our proposed method outperforms previous state-of-the-art method in 14.5% in marco F1 (from 10.3 to 11.8, $P < 0.001$). To further test our model on few-shot setting, we created a new rare diseases coding dataset, MIMIC-III-rare50, on which our model improves marco F1 from 17.1 to 30.4 and micro F1 from 17.2 to 32.6 compared to previous method.

1 Introduction

Multi-label learning has many real-world applications in natural language processing (NLP), including but not limited to academic paper labeling (Chen et al., 2020), news framing (Akyürek et al., 2020), waste crises response (Yang et al., 2020), amazon product labeling (McAuley et al., 2015; Dahiya et al., 2021), and medical coding (Atutxa et al., 2019). In contrast to multi-class classification, an instance in multi-label learning is frequently linked with more than one class label, making the task more challenging due to the combination of potential class labels.

In real-world tasks, there are often insufficient training data for rare class labels. Taking automatic international classification of diseases (ICD) coding as example, given discharge summaries notes as input, the task is to assign multiple ICD disease and procedure label codes associated with each note. The assigned codes need to be accurate and complete for the billing purposes. As an example, the MIMIC-III dataset (Johnson et al., 2016) contains 8,692 unique ICD-9 codes, among which 4,115 (47.3%) codes occur less than 6 times and 203 (2.3%) occur zero times. Clinical practice requires a high accuracy, hence, it is not acceptable for a multi-label classifier to fail a disease diagnosis (or code assignment) because it is rare, since such a diagnosis may be of the most clinical importance for the patient. Therefore, the classifier is required to perform with high precision even for infrequent codes. This translates to data sparsity due to availability of few training examples.

To mitigate the data sparsity problem, additional structured knowledge could be applied. ICD codes are organized with an ontological/hierarchical structure where a text description is associated to each code. For instance, ICD 250 (Diabetes mellitus), shown in Figure 1, is the parent of several child codes including 250.0 (Diabetes mellitus without mention of complication), 250.1 (Diabetes with ketoacidosis), and 250.2 (Diabetes with hyperosmolarity). Such child ICD codes are more semantically different from each other than their parent code 250.

Synonyms including acronyms and abbreviations are common in medical notes. For instance, the description of code 250.00 is disease "type II diabetes mellitus". However, this code can be described in different text forms such as "insulin-resistant diabetes", "non-insulin dependent diabetes", "DM2", and "T2DM". Therefore, one naive way to assign ICD codes is to identify matching between candidate code descriptions and their syn-




Figure 1: An illustration of self-alignment pretraining from medical knowledge UMLS, including the usage of (a) Hierarchy, (b) Synonym, (c) Abbreviation. Pink region is the dynamic margin ranges from $\pi/2$ to π where we wish to pull negatives apart with a dynamic distance.

onyms in medical notes. In this work, we separate synonyms from both acronyms and abbreviations due to their importance in medical domain (Yu et al., 2002). While synonymous relations could be implicitly learned from pretrained language model (LM) (Michalopoulos et al., 2022; Li et al., 2022), previous researches show that language models are only limited to biomedical (Sung et al., 2021) or clinical knowledge bases (Yao et al., 2022) due to the data sparsity challenge in the medical domain. An explicit way of adding such medical knowledge into language model should be explored.

In this paper, we present a simple but effective Knowledge Enhanced PromptT (KEPT) framework. We implement and evaluate KEPT using a LM based on Longformer because clinical notes are typically more than 500 tokens. Specifically, we first pretrain_{mimic} a Longformer LM on MIMIC-III dataset. Then, we further pretrain_{uMLS} on structured medical knowledge UMLS (Unified Medical Language System) using self-alignment learning with contrastive loss to inject medical knowledge into pretrained LM. For the downstream ICD-code assignment fine-tuning, we add a sequence of ICD code descriptions (label semantics) as prompts in addition to each clinical note as KEPT LM input. This allows early fusion of code descriptions and the input note. Experiments on full disease coding (MIMIC-III-full) and common disease coding (MIMIC-III-50) show that our KEPTLongformer outperforms previous SOTA MSMN (Yuan et al.,

2022). In order to test its few-shot ability, we create a new few-shot rare diseases coding dataset named MIMIC-III-rare50, and results show significant improvements compared between MSMN and our method. To facilitate future research, we publicly release the code and trained models¹.

2 Related Work

2.1 Prompt-based Fine-tuning

Prompt-based fine-tuning is effective in few-shot tasks (Le Scao and Rush, 2021; Gao et al., 2021), even when the language model is relatively small (Schick and Schütze, 2021) because they introduce no new parameter during few shot fine-tuning. Additional tuning techniques such as to tune bias-term or language model head have shown to be efficient on memory and training time (Ben Zaken et al., 2022; Logan IV et al., 2022). However, most previous works focus on injecting knowledge into prompt on single-label multi-class classification task (Hu et al., 2022; Wang et al., 2022a; Ye et al., 2022). To the best of our knowledge, this is the first work that applies prompting to multi-label classification task.

2.2 Entity Representation Pretraining

Many recent researches use synonyms to conduct biomedical entity representation learning (Sung et al., 2020; Liu et al., 2021; Lai et al., 2021; Angell et al., 2021; Zhang et al., 2021; Kong et al.,

¹<https://github.com/whaleloops/KEPT>

2021; Seneviratne et al., 2022). Our work is most similar to Liu et al. (2021), which uses additional pretraining scheme that self-aligns the representation space of biomedical entities from pretrained medical LM. They collect self-supervised synonym examples from the biomedical ontology UMLS, and use multi-similarity contrastive loss to keep the representation of similar entities closer to each other, before fine-tuning them to the downstream specific task. However, their work differs from ours in (1) their testing being limited to only medical entity linking tasks and (2) not using hierarchical information, which has been shown to be useful in KRISBERT (Zhang et al., 2021). In contrast to KRISBERT, our contrastive learning selects negative samples from siblings (1-hop nodes) instead of random nodes in the graph. Our method follows InfoMin proposition that selected samples should contain as much task-relevant information while discarding as much irrelevant information in the input as possible (Tian et al., 2020).

2.3 ICD Coding

ICD coding uses NLP models to predict expert labeled ICD codes given discharge summaries as input. Currently, the most straightforward method is to take the best language model for encoding notes, and later use the label attention mechanism to attend labeled ICD codes to input notes for prediction (Mullenbach et al., 2018). In comparison, we apply attention between codes and notes way before within the encoder with the help of prompt. The label representations in attention played an important role in many previous works. Li and Yu (2020) and Vu et al. (2020) first randomly initialize the label representations. Chen and Ren (2019); Dong et al. (2021); Zhou et al. (2021) initialize the label representation with code description from shallow representation using Word2Vec (Mikolov et al., 2013). Yuan et al. (2022) further add description synonyms semantic information. In comparison, we use deep contextual representation from Longformer pretrained on both MIMIC and UMLS with contrastive loss. Similar pretrained language models have shown to be effective in previous works (Wu et al., 2020; Huang et al., 2022; DeYoung et al., 2022; Michalopoulos et al., 2022).

As stated previously, the high dimensions of available label codes, such as 14,000 diagnosis codes and 3,900 procedure codes in ICD-9 and 80,000 in industry coding (Ziletti et al., 2022),

makes ICD coding challenging. Another challenge is the long-tail distribution, in which few codes are frequently used but most codes may only be used a few times due to the rareness of diseases (Shi et al., 2017; Xie et al., 2019). Mottaghi et al. (2020) use active learning with extra human labeling to solve this issue. Other recent works focus on using additional medical domain-specific knowledge to better understand the few training instances (Cao et al., 2020; Song et al., 2020; Lu et al., 2020; Falis et al., 2022; Wang et al., 2022b). Wu et al. (2017) perform entity linking to identify medical phrase in document note. Xie et al. (2019) map label codes as entities in medical hierarchy graph. Compared to a baseline which uses a shallow convolutional neural network to learn n-gram features from notes, they add complex hierarchy structure between codes by allowing the loss to propagate through graph convolutional neural network. In contrast with the previous systems which adopt complex pipelines and different tools, our method applies a much simpler training procedure by incorporating knowledge into language model without requiring any knowledge pre or post-processing (i.e. MedSpacy, Gensim, NLTK) during the fine-tuning. Additionally, previous methods use knowledge graph as an input source, however, we train our language model to include knowledge graph as a target with contrastive loss.

3 Methods

ICD coding: ICD coding is a multi-label multi-class classification task. Specifically, considering thousands of words from an input *medical note* t , the task is to assign a binary label $y_i \in \{0, 1\}$ for each ICD code in the label space Y , where 1 means that note is positive for an ICD disease or procedure and $i \in \text{range}[1, N_c]$. In this study, we define and evaluate the number of candidate codes N_c as 50, although N_c could be higher or lower depending on specific applications. Each candidate code has a short code description phrase c_i in free text. For instance, code 250.1 has description *diabetes with ketoacidosis*. **Code descriptions** c is the set of all N_c number of c_i .

3.1 Encoding Text with Longformer

To solve this task, we first need to encode free text into hidden representation with a pretrained clinical longformer. Specifically, we convert free text a to a sequence of tokens x_a , the vocab embedding then

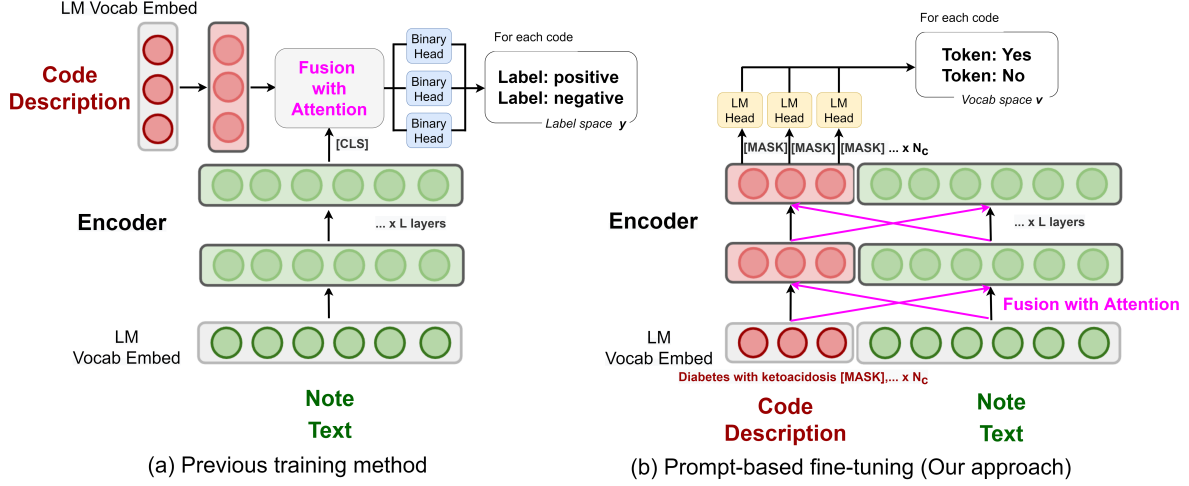


Figure 2: An illustration of (a) standard training method and (b) our proposed prompt-based fine-tuning.

maps x_a to a sequence of hidden vectors. Next, the 1st layer of LM encoder attends one hidden vector to another hidden vector in the sequence with self-attention mechanism. This encoding process is repeated l times to produce a sequence of final contextual hidden vectors $\mathbf{h}_a \in \mathbb{R}^{L_t \times H_d}$ for each free text a where H_d is the hidden layer dimension and L_t is the number of token in t .

3.2 Fine-tuning with Prompt

Prompt based fine-tuning is different from standard fine-tuning. During standard fine-tuning, we usually make input $x_a = [\text{CLS}] a$, where $a \in \{t, c_1, c_2, \dots, c_{N_c}\}$. To assist LM in finding a mention of a label code in note text, we **fusion** final contextual hidden representation of note text t and code description c with attention. Specifically, we first build code description representation $\mathbf{h}'_c \in \mathbb{R}^{N_c \times H_d}$ by concatenating encoded hidden vector $\mathbf{h}_{c_i} \{[\text{CLS}]\} \in \mathbb{R}^{H_d}$ of token [CLS] for each code description c_i . We then build note aware code representation $\mathbf{h}_f \in \mathbb{R}^{N_c \times H_d}$ for each code using cross attention between sequence of vectors \mathbf{h}'_c as query and sequence of vectors \mathbf{h}_t as key, with attention weight α_{ij} between i th item in query and j th item in key as follow:

$$\alpha_{ij} = \text{softmax}((\mathbf{W}_q \mathbf{h}_{c_i} \{[\text{CLS}]\}) (\mathbf{W}_k \mathbf{h}_{t_j}))$$

where \mathbf{W}_q and \mathbf{W}_k are query weight and key weight to be trained. To learn the probability of a code to assign, we train a binary label head, $\text{softmax}(\mathbf{W}_b \mathbf{h}_f)$, by maximizing log-probability of correct label for each code. An illustration of such a standard fine-tuning pipeline is provided in

Figure 2 (a). This standard fine-tuning approach introduces many new parameter weights (589,824 with cross attention and 1,536 with binary label head for longformer), making it hard to learn in few shot settings where the number of training data is limited for each code (Gao et al., 2021). Similar training approaches were carried out in previous researches (Mullenbach et al., 2018; Li and Yu, 2020; Kim and Ganapathi, 2021; Luo et al., 2021; Sun et al., 2021; Zhou et al., 2021) (specific label attention calculation may differ), instead of a pretrained language model, they used unpretrained LSTM or CNN to encode free text, which added more untrained parameters during ICD training.

An alternative approach to multi-label classification is **prompt based fine-tuning**, where masks in prompt are filled-in by LM in cloze style (Gao et al., 2021). We reformulate multi-label classification tasks with free text prompt template as input:

$$x_p = c_1 : [\text{MASK}], c_2 : [\text{MASK}], \dots, c_{N_c} : [\text{MASK}] . t .$$

and use LM to decide if note is positive (or negative) for a code by filling [MASK] with vocab token yes (or no). This step is repeated N_c times for each [MASK] and associated code c_i . Specifically, we encode free text prompt as mentioned before, and obtain final hidden vectors \mathbf{h}_p for input x_p . Notice that this encoding step would **fusion code descriptions** and **note text** with self-attention in every layer of LM encoder. We define a mapping function M from y_i in label space to vocab tokens as:

$$M(y_i) = \begin{cases} \text{"yes"} & \text{if } y_i = 1; \\ \text{"no"} & \text{if } y_i = 0; \end{cases} \quad (1)$$

where $i \in \text{range}[1, N_c]$. In this way, we transfer downstream multi-label classification task into a mask language model task like pretraining. For i th code, the label probability would be calculated as:

$$\begin{aligned} P(y_i|x_p) &= P([\text{MASK}]_{c_i} = M(y_i)|x_p) \\ &= \frac{\exp(\mathbf{W}_{M(y_i)} \cdot \mathbf{h}_p\{[\text{MASK}]_{c_i}\})}{\sum_{\hat{y} \in Y} \exp(\mathbf{W}_{M(\hat{y})} \cdot \mathbf{h}_p\{[\text{MASK}]_{c_i}\})} \end{aligned} \quad (2)$$

where $\mathbf{h}_p\{[\text{MASK}]_{c_i}\} \in \mathbb{R}^{H_d}$ is the hidden vector of the $[\text{MASK}]$ associated with each code c_i in input x_p , and \mathbf{W}_M is the original parameter pre-trained in LM head. Prompt based fine-tuning reuses all parameters during pretraining, and does not introduce new parameters, making the whole model easy to fine-tune in a few-shot setting.

3.3 Hierarchical Self-Alignment Pretrain_{umls} (HSAP) using Knowledge Graph UMLS

Since no new parameters are added to LM in prompt based fine-tuning, the performance on medical downstream task heavily relies on the quality of clinical pretrained LM. However, encoded hidden representations of similar medical terms are not guaranteed to be close to each other. Thus we apply self-alignment pretraining (Liu et al., 2021) to align similar terms closer to each other with additional knowledge. This additional pretrain_{umls} is after masked language pretrain_{mimic} and before auto ICD finetuning. We first build self-supervised data from synonyms, abbreviations, hierarchy in the medical knowledge graph of the UMLS and ICD ontology (§3.3.1), and inject such structural knowledge into a LM by pretraining it on self-supervised data with hierarchical contrastive loss (§3.3.2).

3.3.1 Generating Self-Supervised Data

To generate pretraining examples, we first build a mapping between medical terms and codes as entities in the medical knowledge graph UMLS. Specifically, synonyms of an entity are collected from multiple English free text descriptions of entity via UMLS "MRCONSO" table. Abbreviations of an entity are collected from multiple English free text descriptions of entity in UMLS SPECIALIST Lexicon and Lexical Tools "lraabr" table. Medical terms of an entity are defined as the union of synonyms set and abbreviations set of an entity. To sample negative examples for contrastive loss, we then build a hierarchy tree of entities using ICD-9 code ontology. For example, ICD 250 (Diabetes mellitus) is the parent of ICD 250.0 (Diabetes mellitus without mention of complication), and ICD

250.1 (Diabetes with ketoacidosis) is the sibling of ICD 250.0 (Diabetes mellitus without mention of complication).

3.3.2 Contrastive Learning

Given self-supervised data, we further train clinical longformer using contrastive learning, with the intention of pushing target medical terms and positive medical terms closer, while pulling negative medical terms further away. We formulate this problem into hierarchical triplet loss on a sampled mini batch encoded by LM.

Encoding Medical Terms: Each medical term is usually a short phrase of multiple tokens. Similar to Phrase-BERT (Wang et al., 2021), a medical term is encoded into a sequence of hidden vectors as described in §3.1. We use clinical longformer (Li et al., 2022) as encoder for this process. We define a medical term's hidden representation \mathbf{p} as the first item in hidden vector sequence, which has shown to be effective in Toshniwal et al. (2020).

Hierarchical Neighbor Sampling: We randomly select i number of target anchor entities from the ICD hierarchy level l . Each medical term represents a disease class. Collecting entities from each level could preserve the diversity of samples in the mini batch. Then $j - 1$ parents and siblings are randomly chosen for each of i entities. The purpose of choosing **intra-class** parents and siblings is to encourage model to discriminate anchor entities from close neighbor entities. Finally, k medical terms for each entity are randomly collected, resulting in $n = ijk$ medical terms in a mini batch B of hierarchy level l . We collect mini batch from other hierarchy levels in the same way.

Minibatch Triplet Loss with Dynamic Margin: Similar to Ge et al. (2018), hierarchical triplet loss of a mini batch B can be formulated as:

$$L_B = \frac{1}{2N_B} \sum_{T_x \in T_B} \max(0, m_x - |\mathbf{p}_x^a - \mathbf{p}_x^-| + |\mathbf{p}_x^a - \mathbf{p}_x^+|) \quad (3)$$

where T_B is all the triplets in the minibatch B . N_B is the number of triplets in minibatch B , and each triplet T_x consists of an anchor sample \mathbf{p}_x^a , a positive sample \mathbf{p}_x^+ from positive class, a negative sample \mathbf{p}_x^- from intra-class or inter-class negative class. m_x is a dynamic margin. It is computed according to entity's clinical term similarity between the anchor class entity and the negative class entity (Zakharov et al., 2017). Specifically, for a triplet T_x , the dynamic margin m_x is computed as:

$$m_x = \begin{cases} \pi/2 & \text{if parent;} \\ \pi/2 + \arccos(|\mathbf{p}_x^a \cdot \mathbf{p}_x^-|) & \text{if siblings;} \\ \epsilon & \text{else } (\epsilon = \pi). \end{cases} \quad (4)$$

where condition clause parent and siblings means that negative sample \mathbf{p}_x^- comes from intra-class parent and siblings of anchor sample. In practice, we set $\epsilon = \pi$. Thus, inter-class negative sample would be at least π distance away from anchor sample, while intra-class negative sample would be at least $d \in [\pi/2, \pi]$ range distance away from anchor sample. Such dynamic margin is different from constant margin in previous contrastive loss work in medical domain (Liu et al., 2021; Zhang et al., 2021), and has shown to be effective in visual retrieval task in computer vision (Ge et al., 2018).

By minimizing loss defined in Equation 3, we pretrain a medical knowledge injected clinical longformer. We then use such longformer to encode prompt and context (§3.2), and thus gain knowledge injected prompt for downstream coding task.

When applied to MIMIC-III-full data, it is infeasible to encode all 8,692 candidate ICD codes in the prompt due to high memory cost (to be specified in §6). Instead, we used a two-stages approach. Specifically, we used model MSMN as 1st stage coder to select top 300 candidate codes, and then use our KEPTLongformer as 2nd stage coder to further narrow down the candidates to final prediction. Our 2nd stage coder functions similar to reranker in passage ranking (Nogueira and Cho, 2019).

4 Experiments

4.1 Dataset

MIMIC-III dataset (Johnson et al., 2016) contains data instances of de-identified discharge summaries with expert labeled ICD-9 codes. The discharge summaries are from real patients. We applied the following text pre-processing step before tokenizer: (1) removing all de-identification tokens; (2) replacing characters other than punctuation marks and alphanumeric into white space (e.g. /n); (3) stripping extra white spaces. Previous work (Mullenbach et al., 2018) truncated discharge summaries at 4,000 words. Since longformer used tokens instead of words, we truncated discharge summaries at 8,192 tokens unless otherwise specified. This roughly aligns with our observation that word token ratio is about 1:2. Since procedure codes are related to subjective section of the note (Yang and Yu,

2020), we include relevant sections of discharge summaries for those length exceeds 8,192, and remove irrelevant sections such as discharge followup. The header names of the relevant sections are provided in Table A.1. We named this dataset **MIMIC-III-full**.

For the top-50 frequent codes prediction task, we filtered each instance that has at least one of the top 50 most frequent codes, and used the same splits as the previous work (Vu et al., 2020; Yuan et al., 2022). We named this dataset **MIMIC-III-50**. Detailed statistics are included in Table A.2.

To benchmark auto ICD coding task on few-shot learning, we also created a rare-50 codes prediction using original MIMIC-III dataset. Among 8,692 different types of ICD-9 codes, we first selected codes with less than 10 times occurrences to fit into the few-shot setting. This constitutes more than 90% of original codes. We then ranked the filtered codes by test/train ratio and select top 50, so that testing samples are available for evaluation. We also removed some potential common diseases by hand in the process. This would include true rare diseases (e.g. Kaposi’s sarcoma) listed in expert labeled rare diseases dictionary (Pavan et al., 2017; Wakap et al., 2019). We named this dataset **MIMIC-III-rare50**. The average number of examples per label code (shot) is about 5.

4.2 Implementation Details

For medical domain knowledge graph, we used UMLS 2021AA, containing 4.4 million entities. When mapping an entity to its description, we preferred ICD description. If it is not found, then we used UMLS description. When recreating previous baselines, we used the same hyperparameter setting as mentioned in their published work. We removed R-Drop in (Yuan et al., 2022) and used plain cross-entropy loss only for a fair comparison among all baselines. Code descriptions in prompt use longformer global attention unless otherwise specified. Our full hyperparameter and config setting using wandb is provided in github. Self-alignment Pretraining took about 48 hours with 1 NVIDIA V100 GPU. Fine-tuning took about 10 hours with 2 NVIDIA A100 40GB memory GPUs on MIMIC-III-50, and 0.5 hours on MIMIC-III-rare50. During testing, we used dev set to select best threshold for F1 score. Similar to BERT, no hyper-parameters were further searched on the dev set with our longformer. We evaluated 5 different

Model	AUC		F1		Precision	Best epoch out of 20
	Macro	Micro	Macro	Micro	P@5	
MultiResCNN	89.30	92.04	59.29	66.24	61.56	18
MSATT-KG*	91.40	93.60	63.80	68.40	64.40	-
JointLAAT	92.36	94.24	66.95	70.84	66.36	10
MSMN	92.50	94.39	67.64	71.78	67.23	15
KEPTLongformer	92.63	94.76	68.91	72.85	67.26	4
w/o HSAP	92.33	94.31	67.95	71.92	67.18	5
w/o HSAP & Prompt	90.54	93.18	58.61	67.22	64.38	17
ClinicalBERT	81.94	85.65	43.61	51.62	52.59	15

Table 1: Results on the MIMIC-III-50 test set, compared between KEPTLongformer and baselines (top), KEPTLongformer and ablations (down). * represents result collected from paper because no code is avail.

random seeds for each model and report the median test results across these seeds unless otherwise specified.

4.3 Baselines

MultiResCNN (Li and Yu, 2020) encode free text with Multi-Filter Residual CNN, and applied label code attention mechanism to enable each ICD code to attend different parts of the document.

MSATT-KG (Xie et al., 2019) apply multi-scale attention and graph neural network to capture potential relations between codes, without any changes in the training objectives.

JointLAAT (Vu et al., 2020) propose a hierarchical joint learning with training objectives to predict both ICD code and its parent ICD code in the hierarchy graph.

MSMN (Yuan et al., 2022) use synonyms with adapted multi-head attention, which achieved SOTA performance on MIMIC-III-50 task.

4.4 Results

Results show that our longformer with knowledge pretrained prompt (KEPTLongformer) outperforms the previous state-of-art model MSMN (top of Table 1 and Table 2). For the common disease code assignment (MIMIC-III-50) task, our KEPTLongformer achieves macro AUC of 92.63 (+0.13), micro AUC of 94.76 (+0.36), macro F1 of 68.91 (+1.27), and micro F1 of 72.85 (+1.07). Number in parentheses shows the improvements compared to MSMN. For the rare disease code assignment (MIMIC-III-rare50) task, our KEPTLongformer achieves macro AUC of 82.70 (+7.39), micro AUC of 83.28 (+7.11), macro F1 of 30.44 (+13.39), micro F1 of 32.63 (+15.44). We notice that the im-

provements on rare disease codes are much higher than improvements on common disease code, indicating the strong advantage of our KEPTLongformer for few-shot settings. In contrast to previous work that leads to improvements on rare disease codes but worse results on frequent ones (Rios and Kavuluru, 2018), our approach shows improvements on both tasks. We finally applied our KEPTLongformer to MIMIC-III-full. Table 3 shows that reranker with KEPTLongformer outperforms previous SOTA MSMN in F1 marco from 10.3 to 11.8 by +1.5 (95%CI +0.93 to +1.99, P<0.001) and F1 micro from 58.2 to 59.9 by +1.6 (95%CI +0.95 to +2.33, P<0.001).

4.5 Discussion

Our final KEPTLongformer model could be interpreted as a hybrid of 3 closely interrelated components: longformer, prompt based fine-tuning, and knowledge injected pretraining. Here we provide an ablation study on each part.

Longformer vs. BERT. Increasing max token limit is important under clinical note analysis task, because most clinical notes are long documents with an average of 3000 tokens in MIMIC-III discharge summaries. Due to the high number of tokens in a medical note, it is essential to encode as many tokens as possible before downstream analysis. However, BERT based LM, which could only encode a few sentences, is known to be ineffective for long documents (Beltagy et al., 2020). To test the effect of max token limit in auto ICD coding task, we compare the performance between Clinical Longformer with max limit of 8,192 tokens and ClinicalBert with max limit of 512 tokens. As shown in Table 1, Clinical Longformer

Model	Setting (trained from)	AUC		F1		# Train Param
		Macro	Micro	Macro	Micro	
MSMN	<i>Pretrained</i>	75.3	76.2	17.1	17.2	16.4M
	<i>Finetuned</i>	58.2	44.0	3.3	4.2	16.4M
	<i>Zero shot</i>	52.3	48.9	3.5	4.0	0
KEPTLongformer	<i>Pretrained</i>	81.4	82.3	25.8	30.9	119.4M
	<i>Finetuned</i>	82.7	83.3	30.4	32.6	119.4M
	w/o HSAP	80.2	82.2	24.3	29.9	119.4M
	w/ LM only	75.0	76.9	15.2	16.9	0.6M
	w/ LM & Last	77.6	78.4	17.3	23.4	9.4M
	w/ LM & First	79.0	81.5	23.5	29.6	9.4M
	<i>Zero shot</i>	74.9	76.5	15.2	16.7	0

Table 2: Results on the MIMIC-III-rare50 test set compared between MSMN (previous SOTA on MIMIC-III-50) and our final model KEPTLongformer, where *Pretrained*: model is trained from previous pretraining checkpoint, *Finetuned*: model is trained from best checkpoint after finetuned from MIMIC-III-50, *HSAP*: Hierarchical Self-Alignment Pretraining. We also explore training partial model including: parameters of *LM* head, *Last* self-attention layer, *First* self-attention layer as ablation study. *Zero shot*: No training on rare, directly infer using finetuned model from MIMIC-III-50.

(KEPTLongformer without HSAP & Prompt) substantially outperforms ClinicalBERT in AUC from 7.5 to 8.6 and F1 from 14.9 to 15.6. Other previous methods (e.g. MultiResCNN) use non-pretrained LSTM or CNN with max limit of 8192 tokens. We also observe that these previous methods outperform ClinicalBERT, indicating the importance of max token limit over LM in auto ICD coding task. This finding correlates to previous LM researches (Zhang et al., 2020; Pascual et al., 2021; Biswas et al., 2021) which only uses longformer/BigBird (Michalopoulos et al., 2022) or hierarchical BERT (Ji et al., 2021; Dai et al., 2022) of 4096 max token limits, and our method with max limit of 8192 tokens could alleviate the issues mentioned by them. **Prompt based fine-tuning as early fusion.** In order to test the effect of prompt based fine-tuning as its own, we further compare Longformer trained with prompt based fine-tuning with longformer trained with original fine-tuning on MIMIC-III-50. As shown in Table 1, prompt based fine-tuning (KEPTLongformer w/o HSAP) improves AUC and F1, and converges faster to achieve best F1 score from epoch 17 to epoch 5. Our prompt based finetuned longformer also slightly outperforms MultiResCNN, and other baselines such as MSATT-KG and JointLAAT that uses structured knowledge as addition resources. Under few-shot setting, our prompt based fine-tuning significantly increase AUC and F1 score compared to traditional fine-tuning as shown in Table 2. This finding supports

previous research (Taylor et al., 2022) that shows prompt based fine-tuning outperforms traditional fine-tuning in many few-shot clinical tasks such as length of stay and mortality prediction. Compared to recent models on auto ICD coding, our prompt based model could be seen as an early **fusion** of label code description and input note text. Instead of fusing label description representations and note text representations after encoder with label attention (Zhou et al., 2021; Dong et al., 2021; Yuan et al., 2022), we fuse the two starting from first layer within the encoder with cross attention. Such similar early fusion method has shown to be effective in combining information from knowledge graph and information from text in question answering over knowledge base facts (Das et al., 2017) and open domain question answering (Sun et al., 2018).

Hierarchical self alignment pretraining (HSAP) improves multi-label classification with label domain knowledge. In order to test the effect of HSAP as its own, we further compare Longformer with HSAP (KEPTLongformer) and without HSAP (w/o HSAP). HSAP improves 0.45 on micro AUC and 1.09 on micro F1 in dataset MIMIC-III-50, and 1.1 on micro AUC and 2.7 on micro F1 in dataset MIMIC-III-rare50. Thus we showed that our contrastive learning in label space is more effective in the tasks with limited labeled data, which supports similar finding in text classification (Qian et al., 2022). We also observe that HSAP could reduce

Metric	MSMN	Reranker
F1 Mac	10.3(0.3)	11.8(0.4)
F1 Mic	58.2(0.4)	59.9(0.5)
P@8	74.9(0.3)	77.1(0.3)
R@8	39.2(0.4)	40.7(0.1)
P@15	59.5(0.1)	61.5(0.2)
R@15	55.7(0.1)	57.4(0.2)

Table 3: Results on the MIMIC-III-full compared between previous SOTA MSMN and our final model KEPTLongformer reranker. mean(st.dev.) are reported with 5 different random seeds.

false negative predictions which mistakenly predict their siblings. Out of 78 false negative predictions on code 285.1, 2 predict sibling code 285.9 with HSAP. In contrast, out of 89 false negative predictions on code 285.1, 15 predict sibling code 285.9 without HSAP. HSAP reduces false negative predictions on 285.1 caused by sibling 285.9 from 15 to 2. HSAP works as a good polish to further improve the coding accuracy by injecting domain knowledge into language model.

Parameter efficiency on few-shot learning. One could argue that accuracy improvements come from more number of parameters during training. Our KEPTLongformer is finetuned with 7 times more trainable parameters compared to baseline MSMN. To counter such argument, we also finetune our KEPTLongformer with limited parameters while keeping most parameters fixed. Specifically, we considered the following 4 settings: a) tuning LM head and first encoder layer, b) tuning LM head and last encoder layer, c) tuning LM head only, d) tuning no parameter as zero-shot. Compared to MSMN, settings a, b, c, d improve micro AUC by +5.4, +2.2, +0.7, +0.3 and micro F1 by +12.4, +6.1, -0.3, -0.5 respectively, as shown in Table 2. Setting a and b with 9.4 million trainable parameters significantly outperforms MSMN with 16.4 million trainable parameters. Setting c and d with almost no trainable parameters show competitive results compared to MSMN. We also observe that training first layer outperforms training last layer, this could also be evidence to support the advantage of early fusion for few-shot learning.

5 Conclusions

In this paper, we investigate pretrained clinical language model on auto ICD coding task for both common and rare diseases, the latter of which has

received limited attention in the past. Built on recent advances in contrastive learning, entity representation training and prompt based fine-tuning, our KEPTLongformer easily achieves a competitive performance over state of the art system in common code assignment, and significantly outperforms baseline model in rare code assignment tasks. Finally, our novel Hierarchical Self-Alignment Pre-train could be easily applied to other multi-label classification problems such as tumor detection using other ontology such as OncoTree.

6 Limitations

Our work is limited to auto ICD coding task with 50 label codes including MIMIC-III-50 or MIMIC-III-rare50, and could not be directly applied to ICD coding task MIMIC-III-full with 8,692 labels in practice due to memory constraint. Using our KEPTLongformer would create at least 26,076 tokens and 8,692 [MASK] in a single prompt, which easily explodes the max token limit of a longformer and GPU memory. A more memory efficient method for auto ICD coding could be explored for future work.

Our clinical knowledge pretrained KEPTLongformer is only tested on auto ICD coding task, but such pretrained language model could be easily applied to other clinical NLP applications such as clinical entity linking or clinical question answering tasks. We also only use part of UMLS knowledge graph, including hierarchy, synonym, and abbreviation. Other knowledge including disease co-occurrence, disease-symptom, disease-lab relations and others could also potentially useful for auto ICD coding task.

Acknowledgements

We are grateful to the UMass BioNLP and MLFL group for many helpful discussions and related talks which inspired this work. We would also like to thank the anonymous reviewers for their insightful feedback. Research reported in this study was supported by the National Science Foundation under award 2124126. The work was also in part supported by the National Institutes of Health R01DA045816 and R01MH125027. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Science Foundation and National Institutes of Health.

References

- Afra Feyza Akyürek, Lei Guo, Randa Elanwar, Prakash Ishwar, Margrit Betke, and Derry Tanti Wijaya. 2020. [Multi-label and multilingual news framing analysis](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8614–8624, Online. Association for Computational Linguistics.
- Rico Angell, Nicholas Monath, Sunil Mohan, Nishant Yadav, and Andrew McCallum. 2021. [Clustering-based inference for biomedical entity linking](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2598–2608, Online. Association for Computational Linguistics.
- Aitziber Atutxa, Arantza Díaz de Ilarraza, Koldo Gojenola, Maite Oronoz, and Olatz Perez de Viñaspre. 2019. Interpretable deep learning to map diagnostic texts to icd-10 codes. *International journal of medical informatics*, 129:49–59.
- Iz Beltagy, Matthew E. Peters, and Arman Cohan. 2020. Longformer: The long-document transformer. *ArXiv*, abs/2004.05150.
- Elad Ben Zaken, Yoav Goldberg, and Shauli Ravfogel. 2022. [BitFit: Simple parameter-efficient fine-tuning for transformer-based masked language-models](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 1–9, Dublin, Ireland. Association for Computational Linguistics.
- Biplob Biswas, Thai-Hoang Pham, and Ping Zhang. 2021. Transicd: Transformer based code-wise attention model for explainable icd coding. *ArXiv*, abs/2104.10652.
- Pengfei Cao, Yubo Chen, Kang Liu, Jun Zhao, Shengping Liu, and Weifeng Chong. 2020. [HyperCore: Hyperbolic and co-graph representation for automatic ICD coding](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3105–3114, Online. Association for Computational Linguistics.
- Boli Chen, Xin Huang, Lin Xiao, and Liping Jing. 2020. [Hyperbolic capsule networks for multi-label classification](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3115–3124, Online. Association for Computational Linguistics.
- Yuwen Chen and Jiangtao Ren. 2019. Automatic icd code assignment utilizing textual descriptions and hierarchical structure of icd code. *2019 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pages 348–353.
- K. Dahiya, A. Agarwal, D. Saini, K. Gururaj, J. Jiao, A. Singh, S. Agarwal, P. Kar, and M Varma. 2021. Siamesexml: Siamese networks meet extreme classifiers with 100m labels. In *Proceedings of the International Conference on Machine Learning*.
- Xiang Dai, Ilias Chalkidis, S. Darkner, and Desmond Elliott. 2022. Revisiting transformer-based models for long document classification. *ArXiv*, abs/2204.06683.
- Rajarshi Das, Manzil Zaheer, Siva Reddy, and Andrew McCallum. 2017. [Question answering on knowledge bases and text using universal schema and memory networks](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 358–365, Vancouver, Canada. Association for Computational Linguistics.
- Jay DeYoung, Han-Chin Shing, Luyang Kong, Christopher Winestock, and Chaitanya P. Shivade. 2022. Entity anchored icd coding. *ArXiv*, abs/2208.07444.
- Hang Dong, Víctor Suárez-Paniagua, William Whiteley, and Honghan Wu. 2021. Explainable automated coding of clinical notes using hierarchical label-wise attention networks and label embedding initialisation. *Journal of biomedical informatics*, page 103728.
- Matúš Falis, Hang Dong, Alexandra Birch, and Beatrice Alex. 2022. [Horses to zebras: Ontology-guided data augmentation and synthesis for ICD-9 coding](#). In *Proceedings of the 21st Workshop on Biomedical Language Processing*, pages 389–401, Dublin, Ireland. Association for Computational Linguistics.
- Tianyu Gao, Adam Fisch, and Danqi Chen. 2021. [Making pre-trained language models better few-shot learners](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3816–3830, Online. Association for Computational Linguistics.
- Weifeng Ge, Weilin Huang, Dengke Dong, and Matthew R. Scott. 2018. [Deep metric learning with hierarchical triplet loss](#). In *Proceedings of the European Conference on Computer Vision (ECCV)*.
- Shengding Hu, Ning Ding, Huadong Wang, Zhiyuan Liu, Jingang Wang, Juanzi Li, Wei Wu, and Maosong Sun. 2022. [Knowledgeable prompt-tuning: Incorporating knowledge into prompt verbalizer for text classification](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2225–2240, Dublin, Ireland. Association for Computational Linguistics.
- Chao-Wei Huang, Shang-Chi Tsai, and Yun-Nung Chen. 2022. [PLM-ICD: Automatic ICD coding with pre-trained language models](#). In *Proceedings of the 4th Clinical Natural Language Processing Workshop*, pages 10–20, Seattle, WA. Association for Computational Linguistics.

- Shaoxiong Ji, Matti Holtta, 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. Lehman, Mengling Feng, Mohammad Mahdi Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G. Mark. 2016. MIMIC-III, a freely accessible critical care database. *Scientific Data*, 3.
- Byung-Hak Kim and Varun Ganapathi. 2021. Read, attend, and code: Pushing the limits of medical codes prediction from clinical notes by machines. In *MLHC*.
- Luyang Kong, Christopher Winestock, and Parminder Bhatia. 2021. Zero-shot medical entity retrieval without annotation: Learning from rich knowledge graph semantics. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 2401–2405, Online. Association for Computational Linguistics.
- Tuan Lai, Heng Ji, and ChengXiang Zhai. 2021. BERT might be overkill: A tiny but effective biomedical entity linker based on residual convolutional neural networks. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 1631–1639, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Teven Le Scao and Alexander Rush. 2021. How many data points is a prompt worth? In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2627–2636, Online. Association for Computational Linguistics.
- 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 5:8180–8187.
- Yikuan Li, Ramsey M Wehbe, Faraz S. Ahmad, Hanyin Wang, and Yuan Luo. 2022. Clinical-longformer and clinical-bigbird: Transformers for long clinical sequences. *ArXiv*, abs/2201.11838.
- Fangyu Liu, Ehsan Shareghi, Zaiqiao Meng, Marco Basaldella, and Nigel Collier. 2021. Self-alignment pretraining for biomedical entity representations. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4228–4238, Online. Association for Computational Linguistics.
- Robert Logan IV, Ivana Balazević, Eric Wallace, Fabio Petroni, Sameer Singh, and Sebastian Riedel. 2022. Cutting down on prompts and parameters: Simple few-shot learning with language models. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2824–2835, Dublin, Ireland. Association for Computational Linguistics.
- Jueqing Lu, Lan Du, Ming Liu, and Joanna Dipnall. 2020. Multi-label few/zero-shot learning with knowledge aggregated from multiple label graphs. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2935–2943, Online. Association for Computational Linguistics.
- Junyu Luo, Cao Xiao, Lucas Glass, Jimeng Sun, and Fenglong Ma. 2021. Fusion: Towards automated ICD coding via feature compression. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 2096–2101, Online. Association for Computational Linguistics.
- Julian McAuley, Rahul Pandey, and Jure Leskovec. 2015. Inferring networks of substitutable and complementary products. *Proceedings of the 21st ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- George Michalopoulos, Michal Malyska, Nicola Sahar, Alexander Wong, and Helen Chen. 2022. ICDBig-Bird: A contextual embedding model for ICD code classification. In *Proceedings of the 21st Workshop on Biomedical Language Processing*, pages 330–336, Dublin, Ireland. Association for Computational Linguistics.
- Tomás Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. In *1st International Conference on Learning Representations, ICLR 2013, Scottsdale, Arizona, USA, May 2-4, 2013, Workshop Track Proceedings*.
- Ali Mottaghi, Prathusha Kameswara Sarma, Xavier Atriain, Serena Yeung, and Anitha Kannan. 2020. Medical symptom recognition from patient text: An active learning approach for long-tailed multilabel distributions. *ArXiv*, abs/2011.06874.
- 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.
- Rodrigo Nogueira and Kyunghyun Cho. 2019. Passage re-ranking with bert. *ArXiv*, abs/1901.04085.
- Damian Pascual, Sandro Luck, and Roger Wattenhofer. 2021. Towards BERT-based automatic ICD coding: Limitations and opportunities. In *Proceedings of the 20th Workshop on Biomedical Language Processing*, pages 54–63, Online. Association for Computational Linguistics.
- Sonia Pavan, Kathrin Rommel, María Elena Mateo Marquina, Sophie Höhn, Valérie Lanneau, and Ana Rath. 2017. Clinical practice guidelines for rare diseases: The orphanet database. *PLoS ONE*, 12.

- Tao Qian, Fei Li, Meishan Zhang, Guonian Jin, Ping Fan, and WenHua Dai. 2022. Contrastive learning from label distribution: A case study on text classification. *Neurocomputing*.
- Anthony Rios and Ramakanth Kavuluru. 2018. **Few-shot and zero-shot multi-label learning for structured label spaces**. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3132–3142, Brussels, Belgium. Association for Computational Linguistics.
- Timo Schick and Hinrich Schütze. 2021. **It’s not just size that matters: Small language models are also few-shot learners**. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2339–2352, Online. Association for Computational Linguistics.
- Sandarū Seneviratne, Elena Daskalaki, Artem Lenskiy, and Hanna Suominen. 2022. **m-networks: Adapting the triplet networks for acronym disambiguation**. In *Proceedings of the 4th Clinical Natural Language Processing Workshop*, pages 21–29, Seattle, WA. Association for Computational Linguistics.
- Haoran Shi, Pengtao Xie, Zhiting Hu, Ming Zhang, and Eric P. Xing. 2017. Towards automated icd coding using deep learning. *ArXiv*, abs/1711.04075.
- Congzheng Song, Shanghang Zhang, Najmeh Sadoughi, Pengtao Xie, and Eric Xing. 2020. **Generalized zero-shot text classification for icd coding**. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pages 4018–4024. International Joint Conferences on Artificial Intelligence Organization. Main track.
- Haitian Sun, Bhuwan Dhingra, Manzil Zaheer, Kathryn Mazaitis, Ruslan Salakhutdinov, and William Cohen. 2018. **Open domain question answering using early fusion of knowledge bases and text**. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4231–4242, Brussels, Belgium. Association for Computational Linguistics.
- Wei Sun, Shaoxiong Ji, E. Cambria, and Pekka Marttinen. 2021. Multi-task balanced and recalibrated network for medical code prediction. *ArXiv*, abs/2109.02418.
- Mujeen Sung, Hwisang Jeon, Jinhyuk Lee, and Jaewoo Kang. 2020. **Biomedical entity representations with synonym marginalization**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3641–3650, Online. Association for Computational Linguistics.
- Mujeen Sung, Jinhyuk Lee, Sean Yi, Minji Jeon, Sungdong Kim, and Jaewoo Kang. 2021. **Can language models be biomedical knowledge bases?** In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4723–4734, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Niall Taylor, Yi Zhang, Dan W. Joyce, Alejo J. Nevado-Holgado, and Andrey Kormilitzin. 2022. Clinical prompt learning with frozen language models. *ArXiv*, abs/2205.05535.
- Yonglong Tian, Chen Sun, Ben Poole, Dilip Krishnan, Cordelia Schmid, and Phillip Isola. 2020. What makes for good views for contrastive learning. *ArXiv*, abs/2005.10243.
- Shubham Toshniwal, Haoyue Shi, Bowen Shi, Lingyu Gao, Karen Livescu, and Kevin Gimpel. 2020. **A cross-task analysis of text span representations**. In *Proceedings of the 5th Workshop on Representation Learning for NLP*, pages 166–176, Online. Association for Computational Linguistics.
- Thanh Vu, Dat Quoc Nguyen, and Anthony Nguyen. 2020. **A label attention model for icd coding from clinical text**. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pages 3335–3341. International Joint Conferences on Artificial Intelligence Organization. Main track.
- Stéphanie Nguengang Wakap, Deborah M. Lambert, Annie Olry, Charlotte Rodwell, Charlotte Gueydan, Valérie Lanneau, Daniel Murphy, Yann le Cam, and Ana Rath. 2019. Estimating cumulative point prevalence of rare diseases: analysis of the orphanet database. *European Journal of Human Genetics*, 28:165 – 173.
- Han Wang, Canwen Xu, and Julian McAuley. 2022a. **Automatic multi-label prompting: Simple and interpretable few-shot classification**. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5483–5492, Seattle, United States. Association for Computational Linguistics.
- Shufan Wang, Laure Thompson, and Mohit Iyyer. 2021. **Phrase-BERT: Improved phrase embeddings from BERT with an application to corpus exploration**. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10837–10851, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Tao Wang, Linhai Zhang, Chenchen Ye, Junxi Liu, and Deyu Zhou. 2022b. **A novel framework based on medical concept driven attention for explainable medical code prediction via external knowledge**. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 1407–1416, Dublin, Ireland. Association for Computational Linguistics.
- Honghan Wu, Giulia Toti, Katherine I. Morley, Zina M. Ibrahim, Amos A. Folarin, Richard G. Jackson, Ismail Emre Kartoglu, Asha Agrawal, Clive Stringer, Darren Gale, Genevieve Gorrell, Angus Roberts,

- Matthew T. M. Broadbent, Robert J Stewart, and Richard J. B. Dobson. 2017. Semehr: A general-purpose semantic search system to surface semantic data from clinical notes for tailored care, trial recruitment, and clinical research*. *Journal of the American Medical Informatics Association : JAMIA*, 25:530 – 537.
- Zhuofeng Wu, Sinong Wang, Jiatao Gu, Madian Khabsa, Fei Sun, and Hao Ma. 2020. Clear: Contrastive learning for sentence representation. *ArXiv*, abs/2012.15466.
- Xiancheng Xie, Yun Xiong, Philip S. Yu, and Yangyong Zhu. 2019. Ehr coding with multi-scale feature attention and structured knowledge graph propagation. *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*.
- Qing Yang, Chen Zuo, Xingxing Liu, Zhichao Yang, and Hui Zhou. 2020. Risk response for municipal solid waste crisis using ontology-based reasoning. *International Journal of Environmental Research and Public Health*, 17.
- Zhichao Yang and Hong Yu. 2020. [Generating Accurate Electronic Health Assessment from Medical Graph](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3764–3773, Online. Association for Computational Linguistics.
- Zonghai Yao, Yi Cao, Zhichao Yang, Vijeta Deshpande, and Hong Yu. 2022. Extracting biomedical factual knowledge using pretrained language model and electronic health record context. *AMIA Annual Symposium proceedings. AMIA Symposium, 2022*.
- Hongbin Ye, Ningyu Zhang, Shumin Deng, Xiang Chen, Hui Chen, Feiyu Xiong, Xi Chen, and Huajun Chen. 2022. Ontology-enhanced prompt-tuning for few-shot learning. *Proceedings of the ACM Web Conference 2022*.
- Hong Yu, George Hripcsak, and Carol Friedman. 2002. Mapping abbreviations to full forms in biomedical articles. *Journal of the American Medical Informatics Association*, 9(3):262–272.
- 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.
- Sergey Zakharov, Wadim Kehl, Benjamin Planche, Andreas Hutter, and Slobodan Ilic. 2017. 3d object instance recognition and pose estimation using triplet loss with dynamic margin. *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 552–559.
- Sheng Zhang, Hao Cheng, Shikhar Vashishth, Cliff Wong, Jinfeng Xiao, Xiaodong Liu, Tristan Naumann, Jianfeng Gao, and Hoifung Poon. 2021. [Knowledge-rich self-supervision for biomedical entity linking](#). *CoRR*, abs/2112.07887.
- Zachariah Zhang, Jingshu Liu, and Narges Razavian. 2020. [BERT-XML: Large scale automated ICD coding using BERT pretraining](#). In *Proceedings of the 3rd Clinical Natural Language Processing Workshop*, pages 24–34, Online. Association for Computational Linguistics.
- Tong Zhou, Pengfei Cao, Yubo Chen, Kang Liu, Jun Zhao, Kun Niu, Weifeng Chong, and Shengping Liu. 2021. [Automatic ICD coding via interactive shared representation networks with self-distillation mechanism](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5948–5957, Online. Association for Computational Linguistics.
- Angelo Ziletti, Alan Akbik, Christoph Berns, Thomas Herold, Marion Legler, and Martina Viell. 2022. [Medical coding with biomedical transformer ensembles and zero/few-shot learning](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Industry Track*, pages 176–187, Hybrid: Seattle, Washington + Online. Association for Computational Linguistics.

A Appendix

Section header
chief complaint:
procedure:
history of present illness:
past medical history:
brief hospital course:
discharge diagnosis:
discharge condition:

Table A.1: A list of section header names used to truncate if document token length > 8192.

MIMIC-III-full	train	test
#Doc.	47,723	3,372
Avg #words per Doc.	1,504	1,818
Avg #tokens per Doc.	2,479	3,071
%Doc where #tokens < 512	1.1	0.1
%Doc where #tokens < 4096	89.3	82.3
%Doc where #tokens < 8192	99.6	99.4
Avg #codes per Doc.	15.6	17.9
MIMIC-III-50	train	test
#Doc.	8,066	1,729
Avg #tokens per Doc.	3,008	3,665
%Doc where #tokens < 512	0.5	0.1
%Doc where #tokens < 4096	80.1	67.7
%Doc where #tokens < 8192	97.9	98.7
Avg #codes per Doc.	5.7	6.0
MIMIC-III-rare50	train	test
#Doc.	249	142
Avg #tokens per Doc.	3,462	4,131
%Doc where #tokens < 512	0.1	0.1
%Doc where #tokens < 4096	71.1	55.6
%Doc where #tokens < 8192	96.8	96.5
Avg #codes per Doc.	1.0	1.0

Table A.2: Statistics of MIMIC-III dataset under full codes settings (MIMIC-III-full), 50 common codes settings (MIMIC-III-50), and 50 rare codes settings (MIMIC-III-rare50).

Model	AUC		F1	
	Macro	Micro	Macro	Micro
KEPTLongformer				
InfoNCE	92.48	94.32	68.38	72.01
Hierarchical contrastive loss	92.63	94.76	68.91	72.85

Table A.3: Results on the MIMIC-III-50, compared between KEPTLongformer using hierarchical contrastive loss in this work, compared to InfoNCE which is used in KRISBERT.

Model	AUC		F1		Precision	Best epoch out of 20
	Macro	Micro	Macro	Micro	P@5	
KEPTLongformer						
3 layers	88.65	91.97	57.75	64.67	61.87	18
6 layers	91.99	94.41	66.94	71.33	64.67	7
12 layers	92.63	94.76	68.91	72.85	67.26	4

Table A.4: Results on the MIMIC-III-50, compared between KEPTLongformer using different number of layers.

Model	AUC		F1		Precision	Memory
	Macro	Micro	Macro	Micro	P@5	
KEPTLongformer						
global stride = 1	92.63	94.76	68.91	72.85	67.26	34G
global stride = 3	92.49	94.55	68.43	72.30	67.23	27G
global stride = 5	92.24	94.46	68.09	72.17	66.93	25G
global stride = 10	92.02	94.24	66.98	71.14	65.45	23G

Table A.5: Results on the MIMIC-III-50, compared between KEPTLongformer using different number of global attentions. In prompt code descriptions, we set every *global_stride* number of tokens as (longformer) global attention tokens. For example, global stride = 1 means each token in prompt code descriptions is global attention. *Memory* is the required GPU Memory when *per_device_train_batch_size* = 1.