

A Two-Stage Decoder for Efficient ICD Coding

Thanh-Tung Nguyen[†], Viktor Schlegel[†], Abhinav Kashyap[†], Stefan Winkler^{†¶}

[†]ASUS Intelligent Cloud Services (AICS), Singapore

[¶]Department of Computer Science, National University of Singapore

{thomas_nguyen;victor_schlegel;abhinav_kashyap;stefan_winkler}@asus.com

Abstract

Clinical notes in healthcare facilities are tagged with the International Classification of Diseases (ICD) code; a list of classification codes for medical diagnoses and procedures. ICD coding is a challenging multilabel text classification problem due to noisy clinical document inputs and long-tailed label distribution. Recent automated ICD coding efforts improve performance by encoding medical notes and codes with additional data and knowledge bases. However, most of them do not reflect how human coders generate the code: first, the coders select general code categories and then look for specific subcategories that are relevant to a patient’s condition. Inspired by this, we propose a two-stage decoding mechanism to predict ICD codes. Our model uses the hierarchical properties of the codes to split the prediction into two steps: At first, we predict the parent code and then predict the child code based on the previous prediction. Experiments on the public MIMIC-III data set show that our model performs well in single-model settings without external data or knowledge.

1 Introduction

Medical records and clinical documentation contain critical information about patient care, disease progression, and medical operations. After a patient’s visit, medical coders process them and extract key diagnoses and procedures according to the International Classification of Diseases (ICD) system (WHO, 1948). Such codes are used for predictive modeling of patient care and health status, for insurance claims, billing mechanisms, and other hospital operations (Tsui et al., 2002).

Although the healthcare industry has seen many innovations, many challenges related to manual operations still remain. One of these challenges is manual ICD coding, which requires understanding long and complex medical records with a vast vocabulary and sparse content. Coders must select a

small subset from a continuously expanding set of ICD codes (from around 15,000 codes in ICD 9 to around 140,000 codes in ICD 10 (WHO, 2016)). Therefore, manual ICD coding may result in errors and cause revenue loss or improper allocation of care-related resources. Thus, automated ICD coding has received attention not only from the industry but also from the academic community.

Before the rise of deep learning methods, automated ICD coding methods applied rules or decision tree-based methods (Farkas and Szarvas, 2008; Scheurwegs et al., 2017). The focus has now changed to neural networks using two strands of approaches. The first encodes medical documents using pretrained language models (Li and Yu, 2020; Liu et al., 2021), adapts pretrained language models to make them suitable for the clinical domain (Lewis et al., 2020) or injects language models with medical knowledge such as taxonomy, synonyms, and abbreviations of medical diseases (Yang et al., 2022; Yuan et al., 2022). The second improves the representation of pretrained language models, by capturing the relevance between the document and the label metadata such as their descriptions (Mullenbach et al., 2018; Vu et al., 2020; Kim and Ganapathi, 2021; Zhou et al., 2021), co-occurrences (Cao et al., 2020), hierarchy (Faliss et al., 2019; Vu et al., 2020; Liu et al., 2022), or thesaurus knowledge, such as synonyms (Yuan et al., 2022). Although these approaches are supposed to alleviate problems specific to medical coding such as special vocabulary, a large set of labels, etc., they fall short.

Intuitively, human coders generate the code in two stages: first, the coders select the general codes and then look for specific subcategories that are relevant to a patient’s condition. The advantage of adapting this approach to neural networks is that at each stage of the prediction, we deal with a smaller output space and we can have more confidence when predicting the next stage.

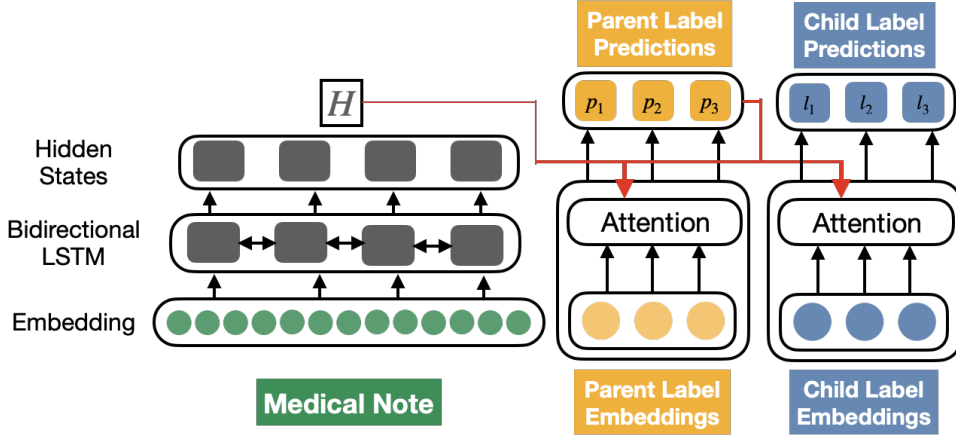


Figure 1: The architecture of our proposed two-stage decoding model. **Encoder (left side)**: the Medical note is embedded and passed through a bidirectional LSTM. **Decoder (right side)**: The first level parent label decoding uses parent label embeddings and attention over the note. The second-level child label predictions are made by attending to the parent label and the note.

Therefore, in this paper, we introduce a simple two-stage decoding framework¹ for ICD coding to mimic the processes of human coders. Our approach leverages the hierarchical structure of the ICD codes to decode, i.e., having parent-child relationships. The first stage predicts the parent codes; the second stage uses the document representation and the predicted parent codes to predict the child codes. Experiments with MIMIC-III data sets demonstrate the effectiveness of our proposed method. In particular, our simple method outperforms models that use external knowledge and data. Since our models (Figure 1) are based on LSTMs, they require less computing power and can be trained faster than other larger models.

2 Two-stage Decoding Framework

ICD codes follow a hierarchical structure. In this work, we consider characters before the dot (.) in the ICD code as the parent label and the code that has to be predicted as the child label. For example, for the child label **39.10** about **Actinomycosis of lung**, its parent code is **39** representing **Actinomyctic infections**. Let \mathbb{P} and \mathbb{L} represent the sets of parent nodes and child codes for a medical note \mathbf{x} , respectively. It is worth noting that if we know the child codes, we can use the above definition to find the corresponding parent codes. This means that knowing \mathbb{L} is equivalent to knowing both \mathbb{L} and \mathbb{P} . Then the probability of the child labels is:

$$\begin{aligned} P_{\theta}(\mathbb{L}|\mathbf{x}) &\sim P_{\theta}(\mathbb{L}, \mathbb{P}|\mathbf{x}) \\ &= P_{\theta}(\mathbb{L}|\mathbb{P}, \mathbf{x})P_{\theta}(\mathbb{P}|\mathbf{x}) \end{aligned} \quad (1)$$

¹<https://github.com/thomasnguyen92/two-stage-decoder-icd>

This factorization allows us to compute the prediction scores of the parent codes first, and then, conditioned on them and the document, we can obtain the prediction score of the child codes. Therefore, we can model the ICD coding task using a decoder framework where we generate parent labels before predicting child labels. In this case, we adapt the decoder framework to the multilabel problem setting, where at each decoding stage, we predict multiple labels at once, instead of one label at a time like a standard decoder.

2.1 Model Architecture

We now describe the components of our parsing model: the document encoder, the first decoding stage for the parent code, and the second decoding stage for the child code.

Document Encoder Given a medical note of n tokens $\mathbf{x} = (x_1, \dots, x_n)$, we embed each token in the document in a dense vector representation. Subsequently, the token representations are passed to a single-layer BI-LSTM encoder to obtain the contextual representations $[h_1, h_2, \dots, h_n]$. Finally, we obtain the encoding matrix $\mathbf{H} \in \mathbb{R}^{n \times d}$.

First Decoding Stage At this stage, similar to [Vu et al. \(2020\)](#), we take the embedding of all parent labels $\mathbf{P} \in \mathbb{R}^{|\mathbb{L}_P| \times d_e}$ to compute the attention scores and obtain the label-specific representations as:

$$\begin{aligned} s(\mathbf{P}, \mathbf{H}) &= \mathbf{P} \tanh(\mathbf{W}\mathbf{H}^T) \\ \text{att}(\mathbf{P}, \mathbf{H}) &= \mathcal{S}(s(\mathbf{P}, \mathbf{H}))\mathbf{H} \\ P(\mathbb{P}|\mathbf{x}) &= \sigma(\text{rds}(\mathbf{V} \odot \text{att}(\mathbf{P}, \mathbf{H}))) \end{aligned}$$

Model	MIMIC-III-Full					MIMIC-III-50				
	AUC		F1		Precision	AUC		F1		Precision
	Macro	Micro	Macro	Micro	P@8	Macro	Micro	Macro	Micro	P@5
Single models										
MultiResCNN (Li and Yu, 2020)	91.0	98.6	9.0	55.2	73.4	89.30	92.04	59.29	66.24	61.56
MSATT-KG (Xie et al., 2019)	91.0	98.6	8.5	55.3	72.8	91.40	93.60	63.80	68.40	64.40
JointLAAT (Vu et al., 2020)	92.1	98.8	10.2	57.5	73.5	92.36	94.24	66.95	70.84	66.36
Our Model	94.6	99.0	10.5	58.4	74.4	92.58	94.52	68.93	71.83	66.72
Models with External Data/Knowledge										
MSMN (Yuan et al., 2022)	95.0	99.2	10.3	58.2	74.9	92.50	94.39	67.64	71.78	67.23
PLM-ICD (Huang et al., 2022)	92.6	98.9	10.4	59.8	77.1	90.23	92.44	65.23	69.26	64.61
KEPTLongformer (Yang et al., 2022)	-	-	11.8	59.9	77.1	92.63	94.76	68.91	72.85	67.26

Table 1: Results on the MIMIC-III-Full and MIMIC-III-50 test sets. All our experiments are run five different random seeds and we report the mean results. The results of other models, except PLM-ICD, are collected from Yang et al. (2022). For PLM-ICD, we follow the authors’ instructions to reproduce the results.

where $\mathcal{S}(\cdot)$, $\sigma(\cdot)$, $\text{rds}(\cdot)$ denote row-wise *softmax*, *sigmoid*, *reduce sum* in last dimension operations; $\mathbf{W} \in \mathbb{R}^{d_e \times d}$ are the weight parameters to perform linear transformations and $\mathbf{V} \in \mathbb{R}^{|L_P| \times d}$ is the weight matrix of a label-wise fully connected layer which yields the parent label logits where \odot is element-wise product.

Second Decoding Stage At this stage, we take the label embeddings of all child labels $\mathbf{L} \in \mathbb{R}^{|L| \times d_e}$ and the probabilities of predicted parent labels from the previous stage as input, and obtain the label-specific representations as per:

$$\begin{aligned}
s(\mathbf{L}, \mathbf{P}) &= \mathbf{L} \tanh(\mathbf{W}_P \odot (P(\mathbb{P}|\mathbf{x}))^T) \\
\text{att}(\mathbf{L}, \mathbf{P}) &= \mathcal{S}(s(\mathbf{L}, \mathbf{P}))\mathbf{P} \\
s(\mathbf{L}, \mathbf{H}) &= \mathbf{L} \tanh(\mathbf{W}_L \mathbf{H}^T) \\
\text{att}(\mathbf{L}, \mathbf{H}) &= \mathcal{S}(s(\mathbf{L}, \mathbf{H}))\mathbf{H} \\
P(\mathbb{L}|\mathbb{P}, \mathbf{x}) &= \sigma(\text{rds}(\mathbf{V}_{LH} \odot \text{att}(\mathbf{L}, \mathbf{H}) \\
&\quad + \mathbf{V}_{LP} \odot \text{att}(\mathbf{L}, \mathbf{P})))
\end{aligned}$$

where we perform a ‘soft’ embedding of the parent labels by taking the element-wise product between matrix $\mathbf{W}_P \in \mathbb{R}^{d_e \times |L_P|}$ with the sigmoid probabilities of parent labels. \mathbf{V}_{LH} , $\mathbf{V}_{LP} \in \mathbb{R}^{|L_P| \times d}$ are the weight matrices of two label-wise fully connected layers that compute the child label logits.

Training Objective & Inference The total training loss is the sum of the binary cross-entropy losses to predict the parent and child labels:

$$\mathcal{L}_{\text{total}}(\theta) = \mathcal{L}_P(\theta) + \mathcal{L}_L(\theta) \quad (2)$$

For inference, we assign a child label to a document if the corresponding parent label score and the child label score are greater than predefined thresholds.

3 Experiments

3.1 Experiment Settings

Setup We conduct experiments on the data set MIMIC-III (Alistair et al., 2016). Following the previous work Joint LAAT (Vu et al., 2020), we consider two versions of MIMIC-III dataset: **MIMIC-III-Full** consisting of the complete set of 8,929 codes and **(MIMIC-III-50)** consisting the 50 most frequent codes. Similarly to Yang et al. (2022), we use macro and micro AUC and F1, as well as precision@k ($k = 8$ for MIMIC-III-Full and $k = 5$ for MIMIC-III-50). For both data sets, we train with one single 16GB Tesla P100 GPU. We detail relevant training hyperparameters and the statistics of the data sets in the Appendix.

We compare our models with recent state-of-the-art work using the results from Yang et al. (2022). Among them, Joint LAAT is most similar to our work because it uses a similar attention mechanism and considers both parent and child labels; therefore, we use it as a comparison in ablation studies. We run our models five times with the same hyperparameters using different random seeds and report the average scores.

3.2 Main Experiment Results

MIMIC-III-Full From the result shown in Table 1, we see that our model achieves a micro F1 of 58.4%, the highest among ‘‘single’’ models that do not rely on external data/knowledge. Specifically, our model outperforms Joint LAAT by about 2.5%, 0.2%, 0.9%, 0.3%, 0.9% in macro AUC, micro AUC, micro F1, macro F1 and precision@8 respectively. In particular, our model is on par with MSMN (Yuan et al., 2022) which uses code synonyms collected from Bodenreider (2004) to




Figure 2: Comparison of Micro-F1 scores between our model and JointLAAT on labels with different MIMIC-III-Full testset frequencies.

improve label embeddings. Moreover, our model is computationally more effective than MSMN with 1.25 vs 11 hours per training epoch and 47 vs. 99 seconds to infer the dev set on a single P100 GPU.

Improvements of other models (Huang et al., 2022; Yang et al., 2022) most likely stem from the use of external information in form of knowledge injected into pre-trained language modeling. We leave the integration of such information into our proposed model architecture for future work.

MIMIC-III-50 From the results on the right-hand side of Table 1, our model produces a micro-F1 of 71.83%, the highest among single models. Specifically, our model surpasses Joint LAAT (Vu et al., 2020) with nearly 1.0%, 2.0%, 0.4% absolute improvement in micro F1, macro F1, and precision@5, respectively. In particular, the macro-F1 of our model is on par with the much more complex state-of-the-art method KEPTLongformer (Yang et al., 2022). This demonstrates the ability of our model to be adapted to classification problems with a large or small number of labels while having competitive results in both cases.

3.3 Ablation Study

To evaluate the effectiveness of our model, we conduct an ablation study on the MIMIC-III-Full set, comparing it with Joint LAAT. Rather than integrating parent label prediction scores as supplementary features with the child label representation, as done in the Joint LAAT method, we allow child label representations to attend to both parent label and document representations. We show that this approach drives performance improvements in two aspects: parent label prediction and performance on labels grouped by frequency of appearance.

Parent Label Prediction Table 2 compares the results of parent label prediction. Our model outperforms Joint LAAT by 0.9%, 0.9%, and 0.7%

	AUC		F1		Precision
	Macro	Micro	Macro	Micro	P@8
Ours	98.7	93.4	29.1	69.0	83.0
JointLAAT	97.8	92.3	28.2	68.1	82.3

Table 2: Parent prediction results on MIMIC-III-Full.

absolute in macro F1, micro F1, and Precision@8, which naturally yields in better child label prediction performance reported in previous sections. But even considering only the case where both models predict parent labels correctly, our approach still achieves a micro F1 score of 65.5%, outperforming Joint LAAT with a micro F1 score of 65.0%. This demonstrates that both parent code and child code prediction benefit from our approach.

Performance in Label Frequency Groups To understand more about our prediction of the model, we divide medical codes into five groups based on their frequencies in MIMIC-III-Full: 1 – 10, 11 – 50, 51 – 100, 101 – 500, > 500 like Wang et al. (2022). We list the statistics of all groups in the Appendix. We compare the micro F1 between different groups in Figure 2. Overall, we outperform Joint LAAT in all groups. The relative improvements are most noticeable in the rare-frequency group (25% relative improvement in the 1 – 10 group, vs 2% or less in other cases). A possible explanation for this is that the parent label space is smaller than the full label space, which results in more training samples per parent label, allowing to learn better representations. As the parent label representation is used to compute child label representations, low-frequency child labels can thus benefit from representations learned from their high-frequency siblings.

4 Conclusion

In this paper, we have presented a novel, simple but effective two-stage decoding model that leverages the hierarchical structure of the ICD codes to decode from parent-level codes to child-level codes. Experiments on the MIMIC-III data set show that our model outperforms other single-model work and achieves on-par results with models using external data/knowledge. Our ablation studies validate the effectiveness of our model in predicting the code hierarchy and codes in different frequency groups. In future work, we intend to integrate our decoder with a better document or label representation to further improve performance.

Limitations

As established, medical coding is an important task for the healthcare industry. Efforts toward its successful automation have wide-ranging implications, from increasing the speed and efficiency of clinical coders while reducing their errors, saving expenses for hospitals, and ultimately improving the quality of care for patients.

However, with this goal in mind, our study presented here should be contextualized by the two main limitations that we identify and outline below.

As with other data-driven approaches, the evaluation performance discussed in our paper is limited by the choice of the (static) MIMIC-III data set. This data set could be seen as lacking diversity, as it only features a fraction of all possible ICD-9 codes and contains medical notes collected in English from a specific department of patients belonging to a specific demographic. While our approach does not make any explicit assumptions about the nature of the data other than the hierarchy of labels, in absence of formal guarantees, we cannot make rigorous statements about the efficacy of our (or indeed any related) approaches on clinical data gathered in different settings, such as other languages, countries or departments.

The second limitation is of a more practical nature, since 2015 the ICD-9 coding system is being phased out in favor of the more expressive ICD-10 system, thus ICD-9 coding has limited applications in practice. However, as with its predecessor, the ICD-10 codes are organized in an even richer hierarchy, which should enable the successful application of our proposed approach.

References

- Johnson Alistair, EW, Pollard Tom, J, and al. et. 2016. MIMIC-III, a freely accessible critical care database. *Scientific data*.
- O. Bodenreider. 2004. The unified medical language system (UMLS): Integrating biomedical terminology. *Nucleic Acids Research*.
- 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.
- Matus Falis, Maciej Pajak, Aneta Lisowska, Patrick Schrempf, Lucas Deckers, Shadia Mikhael, Sotirios Tsaftaris, and Alison O’Neil. 2019. [Ontological attention ensembles for capturing semantic concepts in ICD code prediction from clinical text](#). In *Proceedings of the Tenth International Workshop on Health Text Mining and Information Analysis (LOUHI 2019)*, pages 168–177, Hong Kong. Association for Computational Linguistics.
- Richárd Farkas and György Szarvas. 2008. Automatic construction of rule-based ICD-9-CM coding systems. *BMC Bioinformatics*.
- 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.
- Byung-Hak Kim and Varun Ganapathi. 2021. [Read, attend, and code: Pushing the limits of medical codes prediction from clinical notes by machines](#). In *Proceedings of the 6th Machine Learning for Healthcare Conference*, volume 149 of *Proceedings of Machine Learning Research*, pages 196–208. PMLR.
- Patrick Lewis, Myle Ott, Jingfei Du, and Veselin Stoyanov. 2020. [Pretrained language models for biomedical and clinical tasks: Understanding and extending the state-of-the-art](#). In *Proceedings of the 3rd Clinical Natural Language Processing Workshop*, pages 146–157, Online. Association for Computational Linguistics.
- Fei Li and Hong Yu. 2020. ICD coding from clinical text using multi-filter residual convolutional neural network. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence*.
- Leibo Liu, Oscar Perez-Concha, Anthony Nguyen, Vicki Bennett, and Louisa Jorm. 2022. [Hierarchical label-wise attention transformer model for explainable icd coding](#). *Journal of Biomedical Informatics*, 133:104161.
- Yang Liu, Hua Cheng, Russell Klopfer, Matthew R. Gormley, and Thomas Schaaf. 2021. [Effective convolutional attention network for multi-label clinical document classification](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 5941–5953, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- 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.

Elyne Scheurwegs, Boris Cule, Kim Luyckx, Leon Luyten, and Walter Daelemans. 2017. Selecting relevant features from the electronic health record for clinical code prediction. *Journal of Biomedical Informatics*, pages 74:92–103.

Fu-Chiang Tsui, Michael M. Wagner, Virginia M. Dato, and Chung-Chou Ho Chang. 2002. Value of ICD-9-coded chief complaints for detection of epidemics. *Proceedings. AMIA Symposium*, pages 711–5.

Thanh Vu, Dat Quoc Nguyen, and Anthony Nguyen. 2020. [A label attention model for ICD coding from clinical text](#). In *Proceedings of the 29th International Joint Conference on Artificial Intelligence, IJCAI-20*, pages 3335–3341. Main track.

Tao Wang, Linhai Zhang, Chenchen Ye, Junxi Liu, and Deyu Zhou. 2022. [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.

WHO. 1948. [International statistical classification of diseases and related health problems](#).

WHO. 2016. [International statistical classification of diseases and related health problems, 10th revision](#).

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 28th ACM International Conference on Information and Knowledge Management, CIKM '19*, page 649–658, New York, NY, USA. Association for Computing Machinery.

Zhichao Yang, Shufan Wang, Bhanu Pratap Singh Rawat, Avijit Mitra, and Hong Yu. 2022. Knowledge injected prompt based fine-tuning for multi-label few-shot ICD coding. In *Proceedings of the 2021 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](#). 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.

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.

Appendix

Table 3 depicts the dataset statistic of the MIMIC-III-Full and MIMIC-III-50 datasets, Table 4 details the choice of hyperparameters of our best-performing model and Table 5 outlines the characteristics of the label frequency groupings for the Ablation study in Section 3.3.

	Train	Dev	Test
MIMIC-III Full			
# Doc.	47,723	1,631	3,372
Avg # words per Doc.	1,434	1,724	1,731
Avg # parent codes per Doc.	13.7	15.4	15.9
Total # parent codes	1149	741	850
Avg # child codes per Doc.	15.7	18.0	17.4
Total # child codes	8,692	3,012	4,085
MIMIC-III 50			
# Doc.	8,066	1,573	1,729
Avg # words per Doc.	1,478	1,739	1,763
Avg # parent codes per Doc.	5.3	5.6	5.7
Total # parent codes	39	39	39
Avg # child codes per Doc.	5.7	5.9	6.0
Total # child codes	50	50	50

Table 3: Statistics of MIMIC-III dataset under full codes and top-50 codes settings.

Parameters	Value
Emb. dim. (d_e)	100
LSTM Layer	1
LSTM hidden dim. (h)	512
LSTM output dim. (d)	512
Epoch	50
Batch size	8
Maximum Sequence Length	4000
Optimizer	Adam

Table 4: Hyper-parameters used for training MIMIC-III full and MIMIC-III 50.

Frequency range	Number of parent codes	Number of child codes
1-10	301	5394
11-50	270	1872
51-100	101	549
101-500	246	797
>500	240	309

Table 5: Label Frequency Distribution