

# Read, Attend, and Code: Pushing the Limits of Medical Codes Prediction from Clinical Notes by Machines

**Byung-Hak Kim**  
**Varun Ganapathi**

AKASA

*South San Francisco, CA, USA*

HAK.KIM@AKASA.COM

VARUN@AKASA.COM

**Editor:** Editor's name

## Abstract

Prediction of medical codes from clinical notes is both a practical and essential need for every healthcare delivery organization within current medical systems. Automating annotation will save significant time and excessive effort spent by human coders today. However, the biggest challenge is directly identifying appropriate medical codes out of several thousands of high-dimensional codes from unstructured free-text clinical notes. In the past three years, with Convolutional Neural Networks (CNN) and Long Short-Term Memory (LSTM) networks, there have been vast improvements in tackling the most challenging benchmark of the MIMIC-III-full-label inpatient clinical notes dataset. This progress raises the fundamental question of how far automated machine learning (ML) systems are from human coders' working performance. We assessed the baseline of human coders' performance on the same subsampled testing set. We also present our Read, Attend, and Code (RAC) model for learning the medical code assignment mappings. By connecting convolved embeddings with self-attention and code-title guided attention modules, combined with sentence permutation-based data augmentations and stochastic weight averaging training, RAC establishes a new state of the art (SOTA), considerably outperforming the current best Macro-F1 by 18.7%, and reaches past the human-level coding baseline. This new milestone marks a meaningful step toward fully autonomous medical coding (AMC) in machines reaching parity with human coders' performance in medical code prediction.

## 1. Introduction

Automatic clinical coding (ACC) and emerging AMC are technologies that use natural language processing (NLP) to generate diagnosis and procedure medical codes from clinical notes automatically. A human coder or health care provider scans the medical documentation in electronic health records (EHR), identifying essential information and annotating codes for that particular treatment or service with the current ACC engine. With a wide range of medical services and providers (primary care clinics, specialty clinics, emergency departments, mother-baby units, outpatient and inpatient units, etc.), the complexity of human coders' tasks increases as the medical industry advances, while productivity standards decrease as charts take more time to review. A typical performance for an inpatient coder is 2.5 charts per hour because the review includes capturing historical diagnoses, lab results, radiology results, and several notes from different providers seen during his or her hospital stay.

## Generalizable Insights about Machine Learning in the Context of Healthcare

To better understand human coders’ coding performance baseline, we have hired two professional coders to code the same MIMIC-III inpatient clinical notes testing set that the ML systems were also tested on. We found that the human coders’ inter-agreement rates are not as high as believed, and the RAC model-based ML system exceeds the human coding baseline by a large margin. As far as we know, this is the first paper that benchmarks human coders’ performance on the MIMIC-III dataset by estimating human coders’ ability to agree with one another. Our work makes the three main contributions below:

- First, the human coding baseline is estimated via a primitive internal website. Given the reference diagnosis and procedure codes that Beth Israel Deaconess Medical Center (BIDMC) assigned, we have two coders independently code the different set of notes and evaluate where their total annotations have differed from the references.
- Second, a new RAC architecture that can process unstructured medical notes and attend to text areas annotating medical codes is developed. The main building blocks of RAC are built upon self-attention and code-title guided attention modules that work on sets of sentence vectors. In principle, we look at the medical codes prediction problem as a set-to-set assignment learning problem from the set of input sentence vectors to the set of code labels and employ the problem’s unique permutation equivariant property<sup>1</sup> in the design.
- Third, based on the MIMIC-III dataset, we demonstrate the RAC model’s effectiveness in the most challenging full codes prediction testing set from *inpatient* clinical notes. The RAC model wins over all the previously reported SOTA results considerably. Compared to models that integrate more priors like CNNs & LSTMs, self-attention-based models require more data to generalize well. Hence, to train with the same sized dataset, we utilize the sentence permutation-based data augmentation and stochastic weighted average training (Izmailov et al., 2018).

## 2. Related Work

**Automatic Clinical Coding:** Automatic coding of clinical notes can be treated as a multilabel classification problem. The clinical coding problem has received lots of attention in the last few years from this perspective, and it was demonstrated that CNN and LSTM based deep learning models perform better than conventional machine learning models (Baumel et al., 2017; Shi et al., 2017; Mullenbach et al., 2018; Li et al., 2019; Huang et al., 2019; Umair et al., 2019; Xie et al., 2019; Li and Yu, 2020; Vu et al., 2020); lately, the works have expanded to non-English clinical notes (Azam et al., 2019; Wang et al., 2020; Reys et al., 2020). However, each approach has its disadvantages. The CNN based model requires a stack of many CNN layers to “see” the whole input, and while the LSTM based method is useful at overcoming this, it is not easily parallelizable in training and inference.

**Clinical Transformers:** Converting the free-text in clinical notes into a representation that can easily be used remains one of the prime NLP challenges. With the boom of the

---

1. Consider permuting the notes’ sentence vectors; medical code outputs will be the same but permuted.

Transformer based models, specialized domain versions are trained from scratch or adapted to domains and tasks (see Lee et al., 2020; Peng et al., 2019; Alsentzer et al., 2019; Gu et al., 2020; Zhang et al., 2020). However, successfully applying Transformer models to the medical codes prediction problem remains an open challenge as reported in (Li and Yu, 2020; Ji et al., 2020), primarily due to three reasons: first, free-text clinical notes are unstructured, riddled with spelling errors, and consist of language particular to the medical domain. Second, the output codes space is large, and there is a severe long-tail sparsity issue with over 68,000 codes in the new ICD-10-CM system (CDC, 2015), for example. Third, the standard Transformer-based models’ potential disadvantage in handling longer length of text input than CNN and LSTM based models, which presumably can handle texts with an arbitrary length.

### 3. Human-Level Coding Baselines

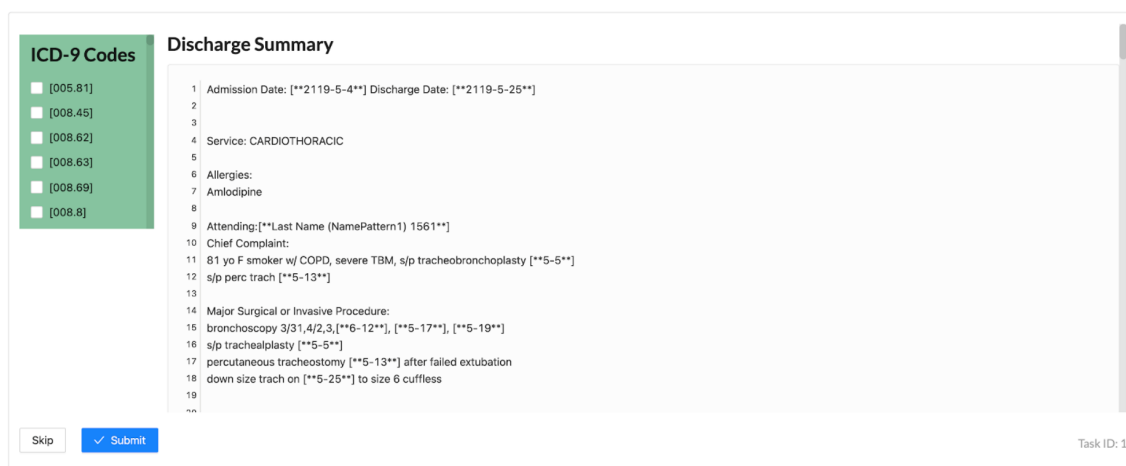


Figure 1: Our web interface for collecting human coders’ assignment of a set of ICD-9 codes to inpatient discharge summaries, described in Section 3. The left column lists the set of ICD-9 codes, and each discharge summary is displayed in the right panel. Note that an entire set of 4,075 codes are too long to show in the screenshot captured. There is a scroll bar next to the ICD-9 codes column that the coder can scroll to click on the appropriate code. Alternatively, the coders can use the find shortcut (e.g., Ctrl + F) and type the code to save time from scrolling to click on the appropriate code. All discharge summaries are sampled from open-access, de-identified MIMIC-III patients’ data treated in intensive care and is not subject to the HIPAA Privacy Rule restrictions on sharing protected health information (Johnson et al., 2016).

**Evaluation Design:** To assess the human coders coding baselines, we hired two professional CPC certified coders with more than five years of hospital coding, including ICD-9 coding experience. We asked them to assign a set of ICD-9 codes (from a shortlisted 4,075

Table 1: Medical code annotation results (in %) by human coders compared to our RAC model-based ML system’s prediction results on 508 random subsamples from the MIMIC-III-full-label testing set. Note that we evaluate the concordance of code assignments between the annotations and the human references in the testing set, as described in Section 3, and the ML system achieves 3.9 times better Micro-Jaccard similarity than human coders.

	<b>Jaccard Similarity</b>		<b>Precision</b>		<b>Recall</b>	
	Macro	Micro	Macro	Micro	Macro	Micro
Human Coders	1.8	<b>10.7</b>	3.5	47.7	2.1	12.1
RAC Model	6.4	<b>42.0</b>	7.1	63.7	10.8	55.2

codes) to a total of 508 inpatient discharge summaries randomly sub-sampled from the MIMIC-III testing set for about 30 hours over a week. The website was built using the Label Studio (Tkachenko et al., 2020) for this task, and each coder is required to log in with their accounts to the website to get started. The coders are given the same discharge summary as was provided to our ML system and were asked to select (not to type) the best possible set of codes in the list for the displayed discharge summaries (see Figure 1 for more details). This environment makes the tasks easier than paper coding, though less convenient than a full-fledged computer-assisted coding system. The intention is to provide the same possible conditions as for the ML system. Additionally, all coders are instructed to code as normally as possible without a time limit to spend on each chart and to use whatever resources they usually use while adhering to coding guidelines. They were asked to code continuously with as few interruptions and distractions as possible, not to communicate with other coders, and not to skip notes.

**Baseline Results:** We calculated the similarity between the human coders we hired and the human coders that the BIDMC hired, available as a human reference in the MIMIC-III dataset, to estimate the human-coders inter-agreement as a human-coding baseline<sup>2</sup>. We also reported how much predicted annotations from our RAC model-based ML system agree with the same MIMIC-III reference to benchmark the ML system performance against the established human-level coding baseline. We summarized the measured human-coders inter-agreement rates in Table 1 in terms of the macro and micro Jaccard similarity, precision, and recall scores and compared them with our RAC-model based automated ML system. We found that estimated inter-coder agreement rates (i.e., human coding baseline) were not high as we initially thought and are far exceeded by the ML system’s 3.9 times higher rate in Micro-Jaccard similarity. We are not here to claim that we reached super-human

2. We measured the human coding baseline by not comparing the agreement of the two certified coders we hired. Instead, we compared the coders we hired with the BIDMC coders. The BIDMC coders’ data is available as a reference label in the MIMIC-III dataset. Of course, neither is an absolute golden truth for inherent ambiguity in the human coding process. Alternatively, we could have two coders independently code, and the notes where they differed are coded independently by a third coder to collect golden reference. We were short on time to pursue this direction when preparing for submission.

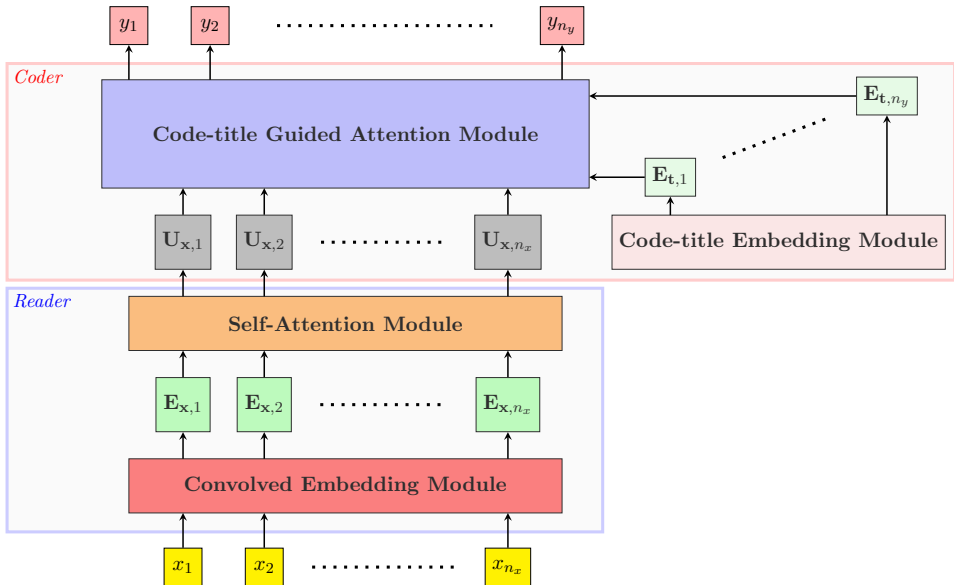


Figure 2: Graphical visualization of Read, Attend, and Code (RAC) model architecture with the reader and coder modules, as described in Section 4.

accuracy as there is room to improve the baseline; instead, this result highlights that our RAC-model based automated ML system is advanced enough to handle complex medical notes to help code more accurately.

## 4. RAC Model

In this section, we describe an end-to-end RAC model to solve the medical code prediction problem. Let  $\mathcal{C}$  be the set of medical codes of size  $n_y$ ,  $\mathcal{C}_T$  be the set of concatenated long and short titles of all  $n_y$  codes,  $\mathbf{x} = (x_1, \dots, x_{n_x})^T$  be the document with  $n_x$  tokens, and  $\mathbf{y} = (y_1, \dots, y_{n_y})^T$  be a medical codes prediction vector. Our RAC model has two sub-modules: a reader and a coder. The reader is a self-attention module that takes a tokenized clinical note  $\mathbf{x}$  as input, and the coder is a code-title guided attention module that predicts each medical code’s likelihood vector  $\mathbf{y}$ . Figure 2 visualizes RAC architecture with these components, and we provide more details in the following sections.

### 4.1. Read

**Convolved Embedding Module:** The *Read* operation first converts each token into a token embedding of dimension  $d$  by using an embedding layer followed by two one-dimensional CNN layers. Unlike prior works (Mullenbach et al., 2018; Xie et al., 2019; Li and Yu, 2020; Ji et al., 2020), we use CNN layers as an efficient preprocessing mechanism to capture an embedding for a group of nearby tokens (like n-grams), which is informative in the next self-attention module for expressing local dependencies in clinical notes. We let  $\mathbf{E}_x \in \mathbb{R}^{n_x \times d}$  be the convolved embedding matrix corresponding to the  $n_x$  tokens in the input document.

Since the next self-attention module does not need to “know” the embeddings’ orders, we do not need to concatenate with positional encoding, unlike the language modelings.

We pre-train the Word2Vec Skip-gram model on all the notes in the training set using Gensim (Řehůřek and Sojka, 2010), with an embedding size of  $d = 300$ , a window size of 5, and a minimum word frequency count of 10 for 5 epochs. The pre-trained weights are then loaded into the embedding layer with a maximum input length of  $n_x = 4096$  tokens. We stack two CNN layers with  $d$  filters, a kernel size of 10, and a tanh activation function on top of the embedding layer. We apply dropout to the module output with a rate of 0.1.

**Self-Attention Module:** In our model, the self-attention module (SAM) is a stack of four identical layers. Each layer is a stack of single-head self-attention and feed-forward layers interleaved with residual connections and layer normalization similar to (Vaswani et al., 2017). In particular, for the convolved embeddings  $\mathbf{E}_x$ , the dot product attention is computed as follows:

$$\text{Attn}(\mathbf{E}_x) = \text{LN} \left( \mathbf{E}_x + \text{Softmax} \left[ \frac{(\mathbf{E}_x \mathbf{W}_q)(\mathbf{E}_x \mathbf{W}_k)^T}{\sqrt{d}} \right] (\mathbf{E}_x \mathbf{W}_v) \right).$$

Here  $\mathbf{W}_q, \mathbf{W}_k, \mathbf{W}_v \in \mathbb{R}^{d \times d}$  represent the projection matrices associated with the query, key, and value, respectively and LN represents the layer normalization. Subsequently, the layer’s output takes the following form:

$$\text{LN} (\text{Attn}(\mathbf{E}_x) + \sigma(\text{Attn}(\mathbf{E}_x) \mathbf{W}_1) \mathbf{W}_2),$$

where  $\mathbf{W}_1 \in \mathbb{R}^{d \times d_{ff}}$ ,  $\mathbf{W}_2 \in \mathbb{R}^{d_{ff} \times d}$ , and  $\sigma$  is a ReLU activation function. We use  $d_{ff} = 1024$  and apply dropout to each sub-layer’s output before it is added to the sub-layer input and normalized with a rate of 0.1.

## 4.2. Attend and Code

**Code-title Embedding Module:** We now illustrate the *AttendAndCode* function to effectively handle the large code output space’s extreme long-tail sparsity. The majority of medical codes are rare, “tail” codes. For example, 60% of ICD codes appear less than 10 times in the MIMIC-III dataset. Moreover, there are roughly 16 codes on average in MIMIC-III discharge summaries, so 99.8% are zeros in  $\mathbf{y}$ . To handle this extreme sparsity of the large code output space, unlike prior sub-optimal attempts (Mullenbach et al., 2018; Wang et al., 2018; Xie et al., 2019), we utilize the code titles (i.e., output labels) information in terms of queries for the next attention module. Code titles corresponding with medical codes are defined by ICD coding systems. Such code titles include, for example, “Intestinal infection due to Clostridium difficile (008.45)” and “Sinoatrial node dysfunction (427.81)”.

We use the definition tables of the diagnoses and procedure codes, concatenate long and short titles together for all  $n_y$  codes, and build  $\mathcal{C}_T$  first. By tokenizing  $\mathcal{C}_T$  with  $n_t$  tokens, we have a title matrix  $\mathbf{T}$  where  $\mathbf{T} \in \mathbb{R}^{n_y \times n_t}$ . From  $\mathbf{T}$  input, the module extracts a code-title embedding of dimension  $d$  by using an embedding layer followed by a single CNN layer and Global Max Pooling layer. We let  $\mathbf{E}_t \in \mathbb{R}^{n_y \times d}$  be the extracted code-title embedding matrix. In the model, each concatenated code title is padded to  $n_t = 36$  tokens, the same pre-trained Word2Vec Skip-gram model weights that the reader used are loaded

to initialize the embedding layer, and a single CNN layer with  $d$  filters, kernel size 10, and tanh activation function are used.

**Code-title Guided Attention Module:** This function computes code-level attention over the reader output to attend to different parts for each code. We explicitly use  $\mathbf{E}_t$  as a query matrix to guide where to attend from the reader output. Specifically, the approach leads to the following attention mechanism:

$$\mathbf{V}_x = \text{Softmax}\left(\frac{\mathbf{E}_t \mathbf{U}_x^T}{\sqrt{d}}\right) \mathbf{U}_x,$$

where  $\mathbf{U}_x \triangleq \text{SAM}(\mathbf{E}_x)$  and  $\mathbf{V}_x \in \mathbb{R}^{n_y \times d}$ .

Fundamental improvement from the per-label attention initially introduced in (Mullenbach et al., 2018) is that  $\mathbf{E}_t$  is no longer randomly initialized. It is straightforward to show that queries close in Euclidean space have similar attention scores, and this property is very desirable for learning inter-relations between less frequent codes and the text effectively. Using  $\mathbf{E}_t$  as a query in computing attention scores has actually resulted in large performance gains, as clearly seen by comparing the first and second row in the RAC Models section of Table 2. This reveals that using learned semantic patterns of code titles in query leads to improving the RAC model’s quality, particularly for the “tail” codes.

With attended  $\mathbf{V}_x$ , finally, the module produces a code likelihood vector  $\mathbf{y}$  as

$$\mathbf{y} = \sigma(\mathbf{V}_x \mathbf{W}_3),$$

where  $\mathbf{W}_3 \in \mathbb{R}^{d \times 1}$  and  $\sigma$  is the sigmoid function.

### 4.3. Learn the Entire Model

**Sentence Permutation:** Unlike computer vision tasks, data augmentation should be done carefully in the NLP application due to the text’s structure. By taking into account the problem’s nature of permutation equivariance, we rely on a simple sentence permutation method. Provided that notes in the training set contain multiple sentences, they are shuffled in a random order to generate a new train sample with the same label. For our model training, we use the 3-fold augmentation to increase the training set size three times and conclude that it is beneficial, as seen by comparing the second and third rows in the RAC Models section of Table 2. One hindsight here is that the second row shows the RAC model’s clear wins over the best prior baselines across the board without providing the augmented data. In other words, this indicates that the performance improvements of the RAC model are not simply the effects of the sentence permutation.

**Stochastic Weighted Averaging (SWA):** With the augmented train data ready, the entire RAC model is trained on the medical codes prediction task to maximize the log-likelihood of  $n_y$  binary classifiers. Rather than using traditional ensembling techniques to combine multiple models to make an averaged final prediction, we apply the SWA approach. Using SWA, we store the running average of model weights during training, so prediction is faster than conventional ensemble methods. In our training, SWA is averaged every 5 epochs from the first epoch, which provides positive gains, as noted by comparing the last two rows in Table 2.

Table 2: Medical codes prediction results (in %) by ML systems on the MIMIC-III-full-label testing set. The bold value shows the best (and highest) SOTA value for each column metric, and underlined numbers indicate the previous SOTA result. The proposed RAC model outperforms all the previously reported SOTA models and achieves new SOTA milestones across all the evaluation metrics. Note that the RAC model has shown an 18.7% relative performance increase in Macro-F1 over the prior best SOTA model. Results for the baseline models are taken from (Mullenbach et al., 2018).

Model	AUC		F1		Precision@n		
	Macro	Micro	Macro	Micro	5	8	15
<b>Baseline Models</b>							
Logistic Regression	56.1	93.7	1.1	27.2		54.2	41.1
SVM				44.1			
CNN	80.6	96.9	4.2	41.9		58.1	44.3
Bi-GRU	82.2	97.1	3.8	41.7		58.5	44.5
<b>CNN-based Models</b>							
CAML (Mullenbach et al., 2018)	89.5	98.6	8.8	53.9		70.9	56.1
DR-CAML (Mullenbach et al., 2018)	89.7	98.5	8.6	52.9		69.0	54.8
MSATT-KG (Xie et al., 2019)	91.0	<b>99.2</b>	9.0	55.3		72.8	58.1
MultiResCNN (Li and Yu, 2020)	91.0	98.6	8.5	55.2		73.4	58.4
<b>LSTM-based Models</b>							
LAAT (Vu et al., 2020)	91.9	<u>98.8</u>	9.9	<u>57.5</u>	<u>81.3</u>	<u>73.8</u>	<u>59.1</u>
JointLAAT (Vu et al., 2020)	<u>92.1</u>	<u>98.8</u>	<u>10.7</u>	<u>57.5</u>	80.6	73.5	59.0
<b>RAC Models (ours)</b>							
RAC - SWA - SentPerm - CodeTitles	92.4	98.9	11.0	57.4	81.5	74.2	58.8
RAC - SWA - SentPerm	<b>94.9</b>	<b>99.2</b>	11.4	58.0	82.3	74.9	59.5
RAC - SWA	94.8	<b>99.2</b>	12.6	58.2	82.6	74.9	59.8
RAC	94.8	<b>99.2</b>	<b>12.7</b>	<b>58.6</b>	<b>82.9</b>	<b>75.4</b>	<b>60.1</b>

## 5. Codes Prediction by Machines

**MIMIC-III Dataset:** The MIMIC-III Dataset (MIMIC v1.4) is a freely accessible medical database containing de-identified medical data of over 40,000 patients staying in the Beth Israel Deaconess Medical Center between 2001 and 2012. For this study, we extract the discharge summaries and the corresponding ICD-9 codes<sup>3</sup>. We perform the same data processing, and data split as stated in (Mullenbach et al., 2018) for a direct comparison

3. There are two reasons that the ICD-9 codes are used in the experimental studies. First, the publicly accessible MIMIC-III dataset used for the study was collected between 2001 and 2012 before the ICD-10 adoption. Second, the MIMIC-III dataset has been used as a standard benchmark in the prior studies, making meaningful head-to-head comparisons with our work. Additionally, we believe the proposed RAC model is not limited to the ICD-9 system. It is anticipated to work for a more complex and sparse environment (like ICD-10 and forthcoming ICD-11 codesets) with the help of employing code-title guided attention. The advantage of employing the RAC model (vs. prior arts) will be even more pronounced.



with prior works. This processing results in 47,724 samples for training, 1,632 and 3,373 samples for validation and testing. One can refer to Table 2 in (Mullenbach et al., 2018) for more dataset statistics.

**Training Details:** For this study, we use an Adam optimizer with a learning rate of  $8e-5$  and pick a batch size of 16 on 4 Nvidia T4 GPUs. All model training runs are carried with the early stopping after the Precision@8 on the validation set does not improve for 3 epochs. The RAC model typically converges within 10 epochs. The model at the epoch of the highest Precision@8 is evaluated on the testing set. Precision@n corresponds to the number of ground truth labels among the n top-scored outputs, and we made this choice because Precision@8 is informative in a production scenario of providing a fixed number of codes to review.

**Performance Results:** We predict a total of 8,921 unique ICD-9 codes (composed of 6,918 diagnosis and 2,003 procedure codes) available in the MIMIC-III-full-label dataset. A smaller subset of code prediction tasks (e.g., TOP-10 or TOP-50 frequent codes) is not considered because it is less challenging and has limited value in our production scenario. We summarized all the previously reported SOTAs values in Table 2 in terms of the macro and micro average AUC, F1, and Precision@n scores to compare with prior works. Note that the RAC model’s number can be further optimized for each metric. Table 2 shows the RAC models’ superior performance over the past SOTA models across the board. What stands out is that the RAC model show very sharp improvements, 18.7% at Macro-F1 and 3.0% at Macro-AUC. These advancements imply that the RAC model learns a better mapping from the clinical text input to the infrequent codes since macro-metrics places more weight on the uncommon codes. Comparing the last four rows in Table 2 highlights each component’s contributions to the overall performance gains in the RAC model.

## 6. Discussion

In this paper, we present for the first time a human-coding baseline for medical code prediction on the subsampled MIMIC-III-full-label inpatient clinical notes testing set task. We have developed an attention-based RAC model that sets the new SOTA records, and the resulting RAC model outperforms the human-coding baseline to a great extent on the same task. The performance improvements can be attributed to effectively learning the common embedding space between the clinical note and medical codes by utilizing attention mechanisms that efficiently address the severe long-tail sparsity issues. This achievement is one step forward to the bigger vision of a fully autonomous ML coding system that autonomously codes the medical charts without input from human coders.

**Limitations:** The current evaluation has the disadvantage of taking only discharge summaries out of the entire inpatient medical chart. We could turn to advanced models to handle more complex inpatient charts containing various service records. We have not discussed how much the RAC model’s accurate prediction performance positively impacts the coding professionals and the healthcare delivery organizations. Conceivably, the computed attention scores can be visualized to inform the notes’ relevant portions to understand the predictions. We intend to evaluate how much this visualized insight along with the presentation of likely medical codes would help coding teams not waste their time and resources as a future study.

**Ethics and Broader Impact:** An automated ML system for medical code prediction, first of all, intends to streamline the medical coding workflow, reduce human coders’ backlog by increasing productivity, and help human coders navigate through extended and complex charts quickly while reducing coding errors (Crawford, 2013). Suppose relevant diagnoses and procedures are annotated with the appropriate codes, human coders can review the record much more quickly and validate the correct codes. Secondly, the automated system intends to reduce the administrative burden on providers, who could instead focus on delivering care rather than learning the nuances of coding. It also helps maximize revenue by capitalizing on the level of specificity noted in the documentation. This will impact many corners of the medical claims billing process by improving coding accuracy and consistency and reducing denials and compliance risk. Moreover, better-automated software can further enhance clinical documentation, making the overall picture of its quality better, eventually redirecting the wasted healthcare costs to more meaningful purposes (Shrank et al., 2019).

## Acknowledgments

During this work, the discussions and actual human coding baseline projects held with Amy Raymond and Grant Messick proved stimulating and helpful. The author also would like to thank Jesse Swidler, Peng Su, and Shan Huang for considerable infrastructure support. The suggestions of Heidi Lim to improve the figures and website visualizations are also gratefully acknowledged.

## References

- Emily Alsentzer, John R. Murphy, Willie Boag, Wei-Hung Weng, Di Jin, Tristan Naumann, and Matthew B. A. McDermott. Publicly available clinical BERT embeddings. In *Clinical Natural Language Processing (ClinicalNLP) Workshop at NAACL 2019*, 2019.
- Sheikh Shams Azam, Manoj Raju, Venkatesh Pagidimarri, and Vamsi Chandra Kasivajjala. Cascadenet: An LSTM based deep learning model for automated ICD-10 coding. In *Advances in Information and Communication (FICC 2019)*, pages 55–74. Springer, 2019.
- Tal Baumel, Jumana Nassour-Kassis, Raphael Cohen, Michael Elhadad, and No’emie Elhadad. Multi-label classification of patient notes a case study on ICD code assignment. *arXiv preprint arXiv:1709.09587*, 2017.
- CDC. Classification of diseases, functioning, and disability. *National Center for Health Statistics; Centers for Disease Control and Prevention*, 2015.
- Mark Crawford. Truth about computer-assisted coding: A consultant, HIM professional, and vendor weigh in on the real CAC impact. *Journal of American Health Information Management Association*, 84(7):24–27, July 2013.
- Yu Gu, Robert Tinn, Hao Cheng, Michael Lucas, Naoto Usuyama, Xiaodong Liu, Tristan Naumann, Jianfeng Gao, and Hoifung Poon. Domain-specific language model pretraining for biomedical natural language processing. *arXiv preprint arXiv:2007.15779*, 2020.

- Jinmiao Huang, Cesar Osorio, and Luke Wicent Sy. An empirical evaluation of deep learning for ICD-9 code assignment using MIMIC-III clinical notes. *Computer Methods and Programs in Biomedicine*, 177:141–153, 2019.
- Pavel Izmailov, Dmitrii Podoprikin, Timur Garipov, Dmitry Vetrov, and Andrew Gordon Wilson. Averaging weights leads to wider optima and better generalization. In *Proceedings of the Thirty-Fourth Conference on Uncertainty in Artificial Intelligence (UAI 2018)*, Monterey, California, USA, 2018.
- Shaoxiong Ji, Erik Cambria, and Pekka Marttinen. Dilated convolutional attention network for medical code assignment from clinical text. In *Proceedings of the 3rd Clinical Natural Language Processing Workshop*, pages 73–78, Online, November 2020. Association for Computational Linguistics.
- Alistair EW Johnson, Tom J Pollard, Lu Shen, Li-wei H Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G Mark. MIMIC-III, a freely accessible critical care database. *Scientific data*, 3:160035, 2016.
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. BioBERT: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*, 36(4):1234–1240, 2020.
- Fei Li and Hong Yu. ICD coding from clinical text using multi-filter residual convolutional neural network. In *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI 2020)*, 2020.
- Min Li, Zhihui Fei, Min Zeng, Fang-Xiang Wu, Yaohang Li, Yi Pan, and Jianxin Wang. Automated ICD-9 coding via a deep learning approach. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 16(4):1193–1202, 2019.
- James Mullenbach, Sarah Wiegrefe, Jon Duke, Jimeng Sun, and Jacob Eisenstein. 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)*, page 1101–1111, New Orleans, Louisiana, 2018. Association for Computational Linguistics.
- Yifan Peng, Shankai Yan, and Zhiyong Lu. Transfer learning in biomedical natural language processing: An evaluation of bert and elmo on ten benchmarking datasets. In *BioNLP 2019: 18th ACL Workshop on Biomedical Natural Language Processing*, 2019.
- Radim Řehůřek and Petr Sojka. Software framework for topic modelling with large corpora. In *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks*, pages 45–50, Valletta, Malta, May 2010. ELRA.
- Arthur D. Reys, Danilo Silva, Daniel Severo, Saulo Pedro, Marcia M. de Souza e Sá, and Guilherme A. C. Salgado. Predicting multiple ICD-10 codes from brazilian-portuguese clinical notes. *arXiv preprint arXiv:2008.01515*, 2020.
- Haoran Shi, Pengtao Xie, Zhiting Hu, Ming Zhang, and Eric P. Xing. Towards automated ICD coding using deep learning. *arXiv preprint arXiv:1711.04075*, 2017.

- William H. Shrank, Teresa L. Rogstad, and Natasha Parekh. Waste in the US health care system: Estimated costs and potential for savings. *JAMA Network*, 322(15):1501–1509, 2019.
- Maxim Tkachenko, Mikhail Malyuk, Nikita Shevchenko, Andrey Holmanyuk, and Nikolai Liubimov. Label Studio: Data labeling software, 2020. URL <https://github.com/heartexlabs/label-studio>. Open source software available from <https://github.com/heartexlabs/label-studio>.
- Abdullah Umair, Jamil Sawar Mohammad, and Ahmed Aftab. A multi-channel convolutional neural network for ICD coding. In *Proceedings of the 2019 IEEE 14th International Conference on Intelligent Systems and Knowledge Engineering (ISKE 2019)*, pages 1178–1184. IEEE, 2019.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems 30 (NeurIPS 2017)*, 2017.
- Thanh Vu, Dat Quoc Nguyen, and Anthony Nguyen. A label attention model for ICD coding from clinical text. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence (IJCAI 2020)*, Monterey, California, USA, 2020.
- Guoyin Wang, Chunyuan Li, Wenlin Wang, Yizhe Zhang, Dinghan Shen, Xinyuan Zhang, Ricardo Henao, and Lawrence Carin. Joint embedding of words and labels for text classification. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2321–2331, Melbourne, Australia, 2018. Association for Computational Linguistics.
- Ssu-ming Wang, Yu-hsuan Chang, Lu-cheng Kuo, Feipei Lai, Yun-nung Chen, Fei-yun Yu, Chih-wei Chen, Zong-wei Li, and Yufang Chungy. Using deep learning for automatic ICD-10 classification from free-text data. *European Journal of Biomedical Informatics*, 16(1), 2020.
- Xiancheng Xie, Yun Xiong, Philip S. Yu, and Yangyong Zhu. 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 2019)*, pages 649–658, Beijing, China, 2019. ACM.
- Zachariah Zhang, Jingshu Liu, and Narges Razavian. BERT-XML: Large scale automated ICD coding using BERT pretraining. In *Proceedings of the 3rd Clinical Natural Language Processing Workshop*, pages 24–34, Online, November 2020. Association for Computational Linguistics.