

Multimodal Machine Learning for Automated ICD Coding

Keyang Xu[♣]

XKY0714@GMAIL.COM

Mike Lam, Jingzhi Pang, Xin Gao, Charlotte Band[♣]

Piyush Mathur, Frank Papay, Ashish K. Khanna, Jacek B. Cywinski, Kamal Maheshwari[△]

Pengtao Xie[♣]

PENGTAO.XIE@PETUUM.COM

Eric P. Xing[♣]

ERIC.XING@PETUUM.COM

[♣]*Petuum Inc, Pittsburgh, PA, USA*

[△]*Cleveland Clinic, Cleveland, Ohio USA*

Abstract

This study presents a multimodal machine learning model to predict ICD-10 diagnostic codes. We developed separate machine learning models that can handle data from different modalities, including unstructured text, semi-structured text and structured tabular data. We further employed an ensemble method to integrate all modality-specific models to generate ICD codes. Key evidence was also extracted to make our prediction more convincing and explainable.

We used the Medical Information Mart for Intensive Care III (MIMIC-III) dataset to validate our approach. For ICD code prediction, our best-performing model (micro-F1 = 0.7633, micro-AUC = 0.9541) significantly outperforms other baseline models including TF-IDF (micro-F1 = 0.6721, micro-AUC = 0.7879) and Text-CNN model (micro-F1 = 0.6569, micro-AUC = 0.9235). For interpretability, our approach achieves a Jaccard Similarity Coefficient (JSC) of 0.1806 on text data and 0.3105 on tabular data, where well-trained physicians achieve 0.2780 and 0.5002 respectively.

1. Introduction

Clinical Relevance The International Classification of Diseases (ICD), endorsed by the World Health Organization (WHO), is a medical classification list of codes for diagnoses and procedures¹. ICD codes have been adopted widely by physicians and other health care providers for reimbursement, storage and retrieval of diagnostic information (Nadathur (2010), Bottle and Aylin (2008)). The process of assigning ICD codes to a patient visit is time-consuming and error-prone. Clinical coders need to extract key information from Electronic Medical Records (EMRs) and assign correct codes based on category, anatomic site, laterality, severity, and etiology (Quan et al. (2005)). The amount of information and complex hierarchy greatly increase the difficulty. Coding errors can further cause billing mistakes, claim denials and underpayment (Adams et al. (2002)). Therefore, an automatic and robust ICD coding process is in great demand.

EMRs store data in different modalities, such as unstructured text, semi-structured text, and structured tabular data (Bates et al. (2003)). Unstructured text data includes notes

1. <https://www.cdc.gov/nchs/icd/icd10cm.htm>

taken by physicians, nursing notes, lab reports, test reports, and discharge summaries. Semi-structured text data refers to a list of structured phrases and unstructured sentences that describe diagnoses written by physicians. Structured tabular data contains prescriptions and clinical measurements such as vital signs, lab test results, and microbiology test results (Lee and Muis, Lipton et al. (2015), Parthiban and Srivatsa (2012)). How to leverage all information from large-volume EMR data is a non-trivial task. Besides, in clinical practice, providing predictions with black-box machine learning models is not convincing for physicians and insurance companies. How to provide evidence related to predicted ICD codes is also an important task.

Technical Significance In this work, we developed an ensemble-based approach that integrates three modality-specific models to fully exploit complementary information from multiple data sources and boost the predictive accuracy of ICD codes. We also explored methods to produce interpretable and explainable predictions. The primary contribution of this study is twofold:

- We applied NLP and multimodal machine learning to predict ICD diagnostic codes, achieving the state-of-the-art accuracy. The complementary nature of multimodal data makes our model more robust and accurate. In addition, we effectively addressed data imbalance issues, which is a very general problem for ICD code prediction.
- We proposed approaches to interpret predictions for unstructured text and structured tables, which make our prediction more explainable and convincing to physicians.

2. Related Work

Many researchers have explored the topic of automatic ICD coding using text data in recent years (Stanfill et al. (2010)). lar (1996) considered it as a single-label classification on patient discharge summaries with multiple classifiers. Kavuluru et al. (2015) treated it as a multi-label classification problem (Tsoumakas and Katakis (2007)), and developed ranking approaches after text feature engineering and selection using EMR data. Koopman et al. (2015) developed a classification system combining Support Vector Machine (SVM) and rule-based methods to identify four high-impact diseases with patient death certificates. It is also common to see unsupervised and semi-supervised strategies applied in predicting ICD codes. Scheurwegs et al. (2017) proposed an unsupervised medical concept extraction approach using unlabeled corpus to frame clinical text into a list of concepts, which helps with ICD code prediction. Deep learning models are also widely used (Shickel et al. (2018)). Duarte et al. (2018) developed a recurrent neural model and Mullenbach et al. (2018) adopted a CNN-based model with per-label attention mechanism to assign ICD codes based on free-text description. Shi et al. (2017) used a neural architecture with Long Short-Term Memory (LSTM) and attention mechanism that takes diagnostic descriptions as input to predict ICD codes. Similarly, Xie and Xing (2018) further explored tree-of-sequences LSTM encoding and adversarial learning to improve the prediction results.

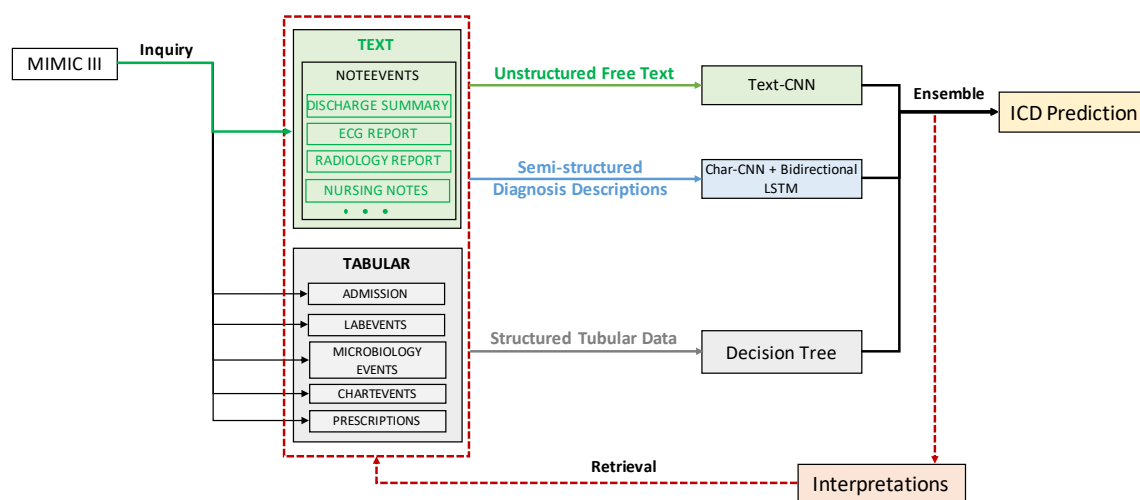


Figure 1: Model architecture for ICD code prediction based on multimodal data, where each prediction is interpreted using retrieval-based methods.

3. Methods

In this section, we will introduce our approach for predicting ICD-10 codes for a patient visit, as well as for identifying the evidence to make our predictions explainable.

We developed an ensemble-based approach as shown in Figure 1, which integrates modality-specific models. For unstructured text, we applied the Text Convolutional Neural Network (Text-CNN) model (Kim (2014)) for multi-label classification. For semi-structured text data, we constructed a deep learning model to analyze semantic similarities between diagnosis written by physicians and ICD code descriptions. For tabular data, our approach transformed numeric features to binary features and fed them to a decision tree (Quinlan (1986)) to classify the ICD codes. During testing, our model ensemble the three aforementioned models on different modalities for improving prediction accuracy (Dietterich (2000)). To make our predictions explainable, key evidence was retrieved from raw data and presented for examination.

4. Results

4.1. Dataset Description

The dataset used for this study is MIMIC-III (Johnson et al. (2016)), which contains approximately 58,000 hospital admissions of 47,000 patients who stayed in the ICU of the Beth Israel Deaconess Medical Center between 2001 and 2012. The original diagnostic codes are ICD-9 codes. We mapped them to ICD-10 codes, as they are more widely adopted in today’s clinical practices. ICD-9 codes with multiple corresponding ICD-10 codes are removed to avoid ambiguity. Based on clinical meaningfulness, 32 ICD codes that are top 50 frequent in both MIMIC-III and a national hospital in the U.S. were selected (see details in Table 3 in Appendix). Patient admissions that have ground truth labels of these 32 codes and at least

one discharge summary were selected. In total we got 44,659 admissions, covering about 77% of the whole dataset. Six major tables from the MIMIC-III dataset were selected:

- ADMISSIONS contains all information regarding a patient admission, including a preliminary diagnosis.
- LABEVENTS contains all laboratory measurements.
- PRESCRIPTIONS contains medications related to order entries.
- MICROBIOLOGYEVENTS contains microbiology information such as whether an organism tested negative or positive in the culture.
- CHARTEVENTS contains all charted data including patients’ routine vital signs and other information related to their health.
- NOTEEVENTS contains all notes including nursing and physician notes, echocardiography reports, and discharge summaries.

4.2. Classification Based on Unstructured Text

In this section, we will discuss how our model predicts diagnostic codes based on unstructured text from NOTEEVENTS. Our approach includes two steps: data pre-processing and deep learning based classification.

The pre-processing step aims to provide a clean and standardized input for the deep learning model. It applied a standardized pipeline including tokenization and word normalization. Tokens with frequency less than 10 were removed. At test time, out-of-vocabulary words were considered to be a special token “UNK”.

The processed text was fed into a multi-label classification model for ICD code prediction. We denote the ICD code set as $\mathcal{Y} = \{1, 2, \dots, C\}$. Given a text input X which is a sequence of tokens, our goal is to select a subset of codes $Y \subseteq \mathcal{Y}$ that is most relevant to X . A Text-CNN model (Kim (2014)) was applied to achieve this. This model represents tokens in X using word embeddings (Mikolov et al. (2013a,b)), and applies a convolution layer on top of the embeddings to capture the temporal correlations among adjacent tokens. To learn the Text-CNN model, the objective is to minimize the average cross-entropy (CE) loss (Shore and Johnson (1980)) for each ICD code class, defined as,

$$\mathcal{L}_{CE} = -\frac{1}{NC} \left[\sum_{i=1}^N \sum_{j=1}^C I_{ij} \log P_{ij} + (1 - I_{ij}) \log (1 - P_{ij}) \right], \quad (1)$$

where N and C are the number of training examples (patients’ admissions) and unique ICD code. I_{ij} is a binary label representing whether example i is assigned with code j and P_{ij} is the probability (predicted by Text-CNN) indicating the likelihood that code j is relevant to example i .

In practice, clinical guidelines developed by human experts are used to guide diagnoses. Medical knowledge from the guidelines can be leveraged by machine learning models to improve prediction accuracy. To achieve it, keywords and phrases were extracted from unstructured guidelines to generate TF-IDF feature vectors (Ramos et al. (2003)). These

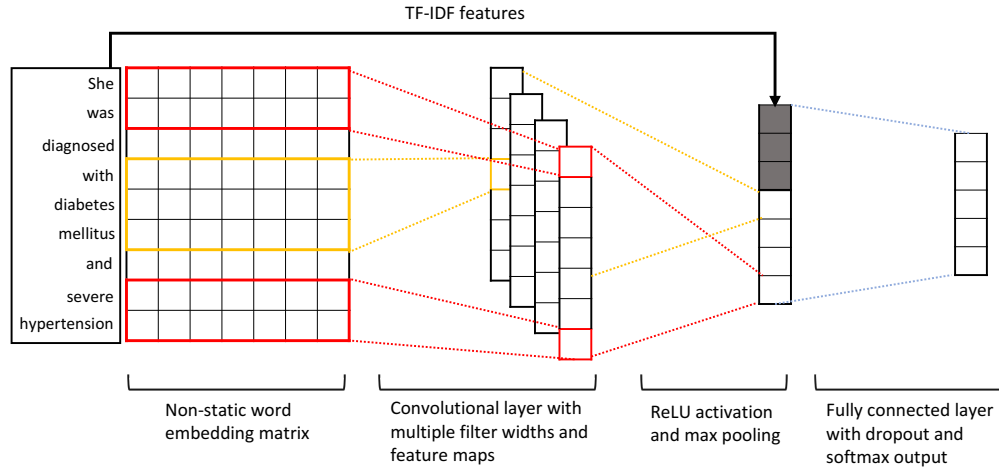


Figure 2: A modified Text-CNN which uses clinical notes as inputs. The model that utilizes TF-IDF features are denoted as Text-TF-IDF-CNN.

features were fed into the fully connected layer of Text-CNN, as additional inputs. We refer to this modified model as Text-TF-IDF-CNN. The detailed network structure is shown in Figure 2.

Training samples for 32 ICD code classes are not evenly distributed – some classes have many samples while others have few. Such data imbalance issue often significantly compromises the performance of machine learning models. To address it, we applied Label Smoothing Regularization (LSR) (Szegedy et al. (2015)). LSR prevents our classifier from being too certain about labels during training and thus avoid overfitting. The LSR approach is able to improve the generalization of our model. More specifically, as defined in (1), $I_{ij} \in \{0, 1\}$ and the average cross-entropy loss is minimized when $P_{ij} = I_{ij}$. LSR was implemented by replacing the truth label I_{ij} with the linear combination between the ground truth label I_{ij} and a sample drawn from a prior distribution $u(j)$, with weights $1 - \epsilon$ and ϵ . Here prior distribution is defined as $u(j) = 1/C$. We sampled ϵ from a beta distribution in each iteration during training, where $\epsilon \sim \text{Beta}(\alpha, \alpha)$ and $\alpha \in (0, \infty)$ is the hyperparameter for Beta distribution. As a result, the smoothed label of class j for sample i is described as,

$$\tilde{I}_{ij} = (1 - \epsilon) \cdot I_{ij} + \epsilon \cdot \frac{1}{C} \quad (2)$$

The average cross-entropy (CE) loss is adjusted in Equation (1) by replacing I_{ij} with the smoothed label \tilde{I}_{ij} , to train our Text-CNN or Text-TF-IDF-CNN model.

4.3. Ranking Based on Semi-structured Text

Medical coders review the diagnosis descriptions taken by physicians in the form of textual phrases and sentences in the clinical notes or admissions table, then manually assign appropriate ICD codes by following the coding guidelines (O’malley et al. (2005)) for billing purpose. For example, the description “Acute on Chronic Kidney Failure” is a strong sig-

nal for the ICD code N17.9 “Acute Kidney Failure, unspecified” because they share close semantic similarity.

We wanted to model this process and formulated this task as Diagnosis-based Ranking (DR) problem where latent representations of all the code descriptions are computed as points in a low-dimensional dense vector space. During inference, diagnosis descriptions are mapped to the same vector space and ICD codes are ranked based on their distance from the diagnosis vector. A neural network model that considers both the character- and word-level features was trained to represent the diagnoses and ICD code descriptions in the same space. To train the model, a triplet loss function is minimized (Cheng et al. (2016)). It takes a triplet as input: *diagnosis* example (\mathbf{x}^a), *positive* code description example \mathbf{x}^p and *negative* code description example \mathbf{x}^n , respectively. The triplet loss defines a relative similarity between instances, by minimizing the distance between the positive pairs and maximizing the distance between the negative ones, defined as:

$$\mathcal{L}_{triplet} = \frac{1}{N} \left[\sum_{i=1}^N \max\{d(\mathbf{x}_i^a, \mathbf{x}_i^p) - d(\mathbf{x}_i^a, \mathbf{x}_i^n) + \alpha, 0\} \right], \quad (3)$$

where $d(\cdot)$ is the Euclidean distance $d(\mathbf{x}, \mathbf{y}) = \|\mathbf{x} - \mathbf{y}\|_2$ for $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$ and α is the margin. Here we took margin α as 1. The associated network structure is shown in Figure 3. Pre-trained word embedding and Character-level Convolution Neural Network (Char-CNN) (Zhang et al. (2015)) are used to represent each word token. The Char-CNN layer creates representations such that tokens with similar form like “injection” and “injections” are mapped closer to each other. Similarly, the word embeddings represent a semantic space, such that a token “Hypertension” and its abbreviation “HTN” map closer to each other. The word embedding matrix was pre-trained on PubMed² (Mikolov et al. (2013b)), which contains abstracts of over 550,000 biomedical papers. Then, a bidirectional LSTM (Huang et al. (2015)) layer encodes the sequential context information followed by a max-pooling layer to compute final representations.

The MIMIC-III dataset does not have a one-to-one mapping of ICD code and diagnosis. To construct the triplets for training, we crawled ICD-10 codes and synonyms for their corresponding code description from online resources³. Then, we mapped all the synonyms for each ICD code as the positive examples. To create the high-quality negative examples, instead of randomly sampling, we computed the edit distance to extract n-grams that are similar to the code description. During inference, min-max normalization transformed the distance value to a range [0, 1]. If one admission has multiple diagnoses, the mean of features was computed to combine those normalized results.

4.4. Classification based on Tabular Data

Tabular data mainly includes four tables, including LABEVENTS, PRESCRIPTIONS, MICROBIOLOGYEVENTS, and CHARTEVENTS.

- LABEVENTS (LAB): It uses a binary value to indicate whether each laboratory value is abnormal. Furthermore, if a patient has multiple records of the same lab test, a

2. <https://www.ncbi.nlm.nih.gov/pubmed/>

3. <https://www.icd10data.com/>

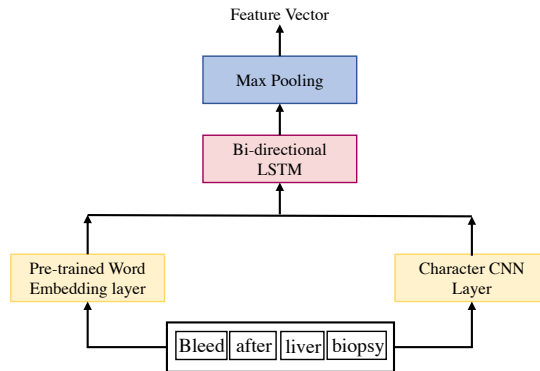


Figure 3: Model architecture for diagnosis-based ranking.

majority vote strategy was adopted. Missing values were considered as normal. In our experiments, a total of 753 lab tests were used to construct features for classification.

- CHARTEVENTS: Based on physicians’ suggestion, we mainly focused on three measurements: Body Mass Index (BMI), heart rate, and blood pressure. A binary vector was used to denote whether a specific measurement result is within the normal range. We combined features from this table with LABEVENTS.
- PRESCRIPTIONS (MED): A binary vector was used to denote whether a medication is prescribed. Missing medications were considered as not being prescribed. Medications with a frequency less than 50 were removed to reduce noise. A total of 1135 medications were used to construct binary features.
- MICROBIOLOGYEVENTS (BIO): A binary feature was used to represent a microbiology even — whether an organism test is positive or negative. A total of 363 organisms were used to form feature vectors.

Our solution applied a decision tree (Quinlan (1986)) as the classifier using binary features and leveraged a one-versus-all strategy (Rifkin and Klautau (2004)) for multi-label classification. Higher weights were given to samples from minority classes to handle data imbalance.

4.5. Model Ensemble

During testing, our solution takes an ensemble of the trained models for predicting ICD codes. Specifically, for each class j , the final predicted probability P_j is a weighted sum of probabilities predicted by individual models:

$$P_j = \sum_{k=1}^K \alpha_k P_j^{(k)}, \tag{4}$$

where $P_j^{(k)}$ is the probability predicted by model k on class j , and K is the total number of models. The weight parameters $\{\alpha_k\}_{k=1}^K$ were tuned by performing grid search on the validation set. However, not every single patient visit is guaranteed to have all the data

resources available. Hence in order to ensure all weights sum up to one, i.e., $\sum_{k=1}^K \alpha_k = 1$, if the k -th predictor is missing, its weight α_k will be given to the Text-CNN or Text-TF-IDF-CNN model.

4.6. Methods for Interpretation

Two separate methods were adopted to extract the important textual and tabular features that have the highest influence on the final ICD code predictions.

4.6.1. TEXT INTERPRETABILITY

After our model predicts a specific ICD code $y \in \{1, 2, \dots, C\}$, it is desirable to identify key phrases that lead to such a prediction from the textual input. To capture the association between a word w and an ICD code y , we extracted all the paths connecting w and y from the trained neural network. For each path, we computed the influence score by multiplying the values of all hidden units and the weights associated with all edges along this path. The scores of all paths were added up to measure the association strength between w and y . Consecutive words with non-zero scores were combined into phrases and then ranked by the maximum score. Top ranked phrases were considered important signals for models to determine a specific ICD code.

4.6.2. TABLE INTERPRETABILITY

The nature of structured tabular features makes the method inherently different from those using text data. Therefore, the Local Interpretable Model-Agnostic Explanation (LIME) method (Ribeiro et al. (2016)) was adopted to tackle this problem. LIME computes a score for each feature of an instance that represents the importance of this feature in contributing to the model’s final prediction. Instead of going through the trained model, LIME learns the weights of features by approximating the decision boundary locally (i.e. instance-specific) based on a simple model.

5. Experiments

In this section, we will introduce experiments for validating our approach and evaluation results for both classification and interpretability.

5.1. Experimental Settings

The data was randomly split into training, validation, and test sets containing 31,155, 4,484, and 9,020 admissions respectively. Admissions of the same patient were categorized into the same set. This prevents the model from memorizing information of a patient from the training set and leveraging that information to inflate performance on the test set.

5.2. Hyperparameters

Hyperparameters were tuned on the validation set and then applied to the test set, including: (1) For Text-CNN and Text-TF-IDF-CNN model, kernel sizes of 2, 3, 4 were adopted and each kernel had 128 feature maps. Word embedding dimension was 256. Dropout rate was

0.1. L2-regularization coefficient was 10^{-4} . Adam optimizer (Kingma and Ba (2014)) with learning rate 10^{-3} was used. Batch size was 32; α was 0.3 for beta distribution in label smoothing regularization. (2) For diagnosis ranking, character embedding size was 50. One layer of bi-directional LSTM with 100 hidden units was used for encoding.

5.3. Evaluation Metrics

- **For Classification**, we used F1 and the Area Under the ROC Curve (AUC) (Goutte and Gaussier (2005); Huang and Ling (2005)) as evaluation metrics. F1 score is the harmonic mean of precision and recall, and AUC score summarizes performances under different thresholds. To better compute the average across different classes, we adopted micro-averages and macro-averages. Classes with more samples have larger weights for micro-averaged metrics but are treated equally for macro-averaged metrics.
- **For Interpretability**, we used the Jaccard Similarity Coefficient (JSC) (Niwattanakul et al. (2013)) to measure the overlap between two sets, which are our extracted evidence and physicians’ annotations. It is defined as,

$$J(A, B) = \frac{|A \cap B|}{|A \cup B|}$$

5.4. Evaluation of Code Classification

The overall performance for multi-label classification is shown in Table 1. Word-TF-IDF and MetaMap-TF-IDF are baseline methods, where MetaMap extracts medical entities and links them to UMLS concepts (Aronson (2001)). Keyword-TF-IDF only adopts TF-IDF values of keywords from clinical guidelines. We used gradient boosting decision tree (Stanford (2002)) model as the classifier.

Most models in Table 1 were developed based on the CNN architecture. Vanilla Text-CNN model and DenseNet (Huang et al. (2017)) perform similarly with TF-IDF models on F1 but achieve better results on AUC.

Label Smoothing (LS) significantly improves the performance as it alleviates data imbalance issue, especially for F1 scores. TextCNN+DR also enhances Text-CNN performance on all four metrics.

For Tabular Data (TD) that comprises BIO, LAB and MED, we observe improvements on all four metrics after ensembling any type with Text-CNN+LS. Incorporating TD, Text-CNN+LS+TD gets the best macro- and micro-F1 (0.6137 and 0.6929 respectively) among all models with Text-CNN structure. By further ensembling DR, our solution gets a higher AUC score but slightly lower F1 scores. Text-CNN+LS+DR+TD improves 13% on macro-F1 and 5.4% on micro-F1 over Text-CNN, and 25.9% on macro-F1 and 19.5% on micro-F1 over Word-TF-IDF, which demonstrates the effectiveness of our proposed ensemble approach and the benefits of using multimodal data.

In addition, evaluation results show that the use of the most relevant subsets from clinical guidelines, denoted as Text-TF-IDF-CNN, can significantly increase the performance. The combined model of Text-TF-IDF-CNN, LS, DR and TD achieves 0.6867 Macro-F1, 0.7633 Micro-F1, 0.9337 Macro-AUC and 0.9541 Micro-AUC, the highest scores among all model tested. Detailed evaluation results for 32 ICD classes are listed in Figure 4.

Table 1: Evaluations of classification for our models and baselines. LS stands for label smoothing. DR is the diagnosis-based ranking model. LAB contains both LABEVENTS and CHARTEVENTS features. Tabular Data (TD) includes Bio, Med and Lab.

Methods	F1		AUC	
	Macro	Micro	Macro	Micro
Word-TF-IDF	.5277	.6648	.7249	.7834
Word-TF-IDF + MetaMap-TF-IDF	.5336	.6721	.7298	.7879
Keyword-TF-IDF	.5678	.7225	.7375	.7867
DenseNet	.5327	.6621	.8885	.9228
Text-CNN	.5429	.6569	.8959	.9235
Text-CNN + LS	.6054	.6908	.9010	.9293
Text-CNN + LS + DR	.6097	.6914	.9029	.9307
Text-CNN + LS + Bio	.6108	.6922	.9076	.9323
Text-CNN + LS + Lab	.6107	.6925	.9136	.9366
Text-CNN + LS + Med	.6122	.6928	.9124	.9379
Text-CNN + LS + TD	.6137	.6929	.9177	.9406
Text-CNN + LS + DR + TD	.6133	.6921	.9188	.9416
Text-TF-IDF-CNN + LS	.6813	.7632	.9157	.9420
Text-TF-IDF-CNN + LS + DR + TD	.6867	.7633	.9337	.9541

5.5. Interpretability Evaluation

We collected a test set of 25 samples from five selected ICD-10 codes, including I10, I50.9, N17.9, E11.9, and D64.9. Annotations were obtained independently from three practicing physicians, who were asked to annotate key evidence from all data sources that determines the assignment of ICD-10 codes. The physicians have 15, 22 and 32 years of experience, respectively.

5.5.1. TEXT-ORIENTED EVALUATION

We compare top- k phrases extracted by our model with physicians’ annotations. k is the number of snippets physician annotated. If tokenized phrases and annotations coincide above a certain threshold, they are considered overlapped. This also applies when comparing overlap between physician’s annotations. We find different physicians have different annotation styles. For instance, physicians 2 and 3 tend to highlight keywords while physician 1 tends to highlight sentences, resulting in fewer tokens obtained from physicians 2 and 3. Therefore, average overlap score is as low as 0.2784 among highly-trained professionals, indicating that finding evidence from clinical notes to determine ICD codes is a non-trivial task.

Table 2 shows the overlap score between physicians’ annotations and the outputs of our model. On average, our model obtains JSC of 0.1806. Our approach is designed to find n-

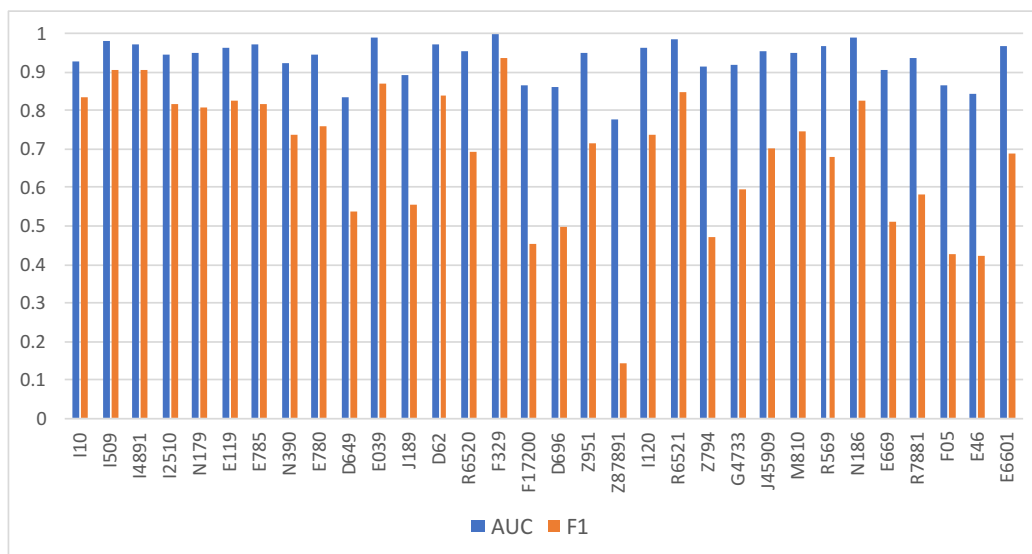


Figure 4: Detailed AUC and F1 scores for 32 ICD codes for Text-TF-IDF-CNN + LS + TD + DR model. ICD codes are listed on the X axis by the training sample size, in descending order.

Table 2: Jaccard Similarity Coefficient (JSC) between physicians and our model

	JSC	Physician 1	Physician 2	Physician 3	Our Model
Text Data	Physician 1	-	.3105	.1798	.1477
	Physician 2	.3105	-	.3450	.2137
	Physician 3	.1798	.3450	-	.1905
Tabular Data	Physician 1	-	.5455	.4225	.3867
	Physician 2	.5455	-	.5325	.2885
	Physician 3	.4225	.5325	-	.2564

grams carrying important signals. It can capture numerous phrases that are either directly related to a specific disease or that can provide insights for the final code prediction.

5.5.2. TABLE-ORIENTED EVALUATION

For each table, the k most important features found by LIME were selected as evidence for each sample, where k is the number of important features a physician annotated. We removed duplicates and computed the JSC. All results are shown in Table 2. We observe that the average agreement among physicians on tabular data is 0.5002, higher than that on text data. The average JSC between our model and physicians’ annotations is 0.3105. In general, our model captures more features than physicians. Some of those extra features are informative for diagnoses even though they are not annotated by physicians. This is similar to the finding from text interpretability.

6. Discussion

6.1. Code Classification

The use of TF-IDF features makes vector space model perform well on F1 score, indicating the importance of retrieving disease-related keywords in these classification tasks. Inspired by Keyword-TF-IDF that filters out noise and irrelevant features, we proposed Text-TF-IDF-CNN that incorporates cognitive TF-IDF features extracted from clinical guidelines to the fully-connected layer in our deep learning model. It shows the possibility of leveraging human domain knowledge to improve ML model performance in clinical settings.

Performance improvements on macro-averaged metrics are greater than micro-averaged metrics. One reason is that multimodal data mitigates the issue of limited inference from text data caused by small training sample size. Additionally, given physical examination and lab tests are mandatory for most inpatients, tabular data considerably benefits classes with a small sample size. As we expand the prediction scope to cover larger code list in the future, resolving data imbalance will play a bigger role in improving the predictive accuracy.

6.2. Interpretability

For text data, key evidence was extracted based on the Text-TF-IDF-CNN model. For tabular data, key evidence was extracted based on LIME from LABEVENTS and PRESCRIPTIONS tables.

For text data, among all ranked evidence extracted, our model is able to identify keywords highly relevant to a specific code. For example, “DM II” and “Metformin” were extracted for E11.9. Results reveal that our model is able to identify text snippets that are related to a specific disease, even though the relationship can be indirect. Take N17.9 for example, the snippet “IVDU, states he used heroin, barbiturates, cocaine” indicates drug abuse could be the cause of acute kidney failure.

For tabular data, valuable items can be extracted as well, especially where text evidence is not sufficient. For example, “Glucose” from LABEVENTS and “Insulin” from PRESCRIPTIONS were extracted as top evidence for E11.9, which were not found in top ranked keywords from unstructured text.

Interpretation results are affected by the performance of classification. For example, items extracted in classes with an AUC score above 0.9 are more relevant than those in class D64.9 with an AUC score around 0.7 (see Figure 4).

Table 4 in Appendix shows some examples of evidence extracted from multiple modalities.

6.3. Future Work

We plan to improve our models in following areas:

- Enlarge the code list. Currently we focus on 32 ICD-10 codes that are representative in hospitals and we will cover more minority classes.
- Reduce feature dimensions. Currently feature dimensions for tabular data are very large and some features are duplicate. Feature dimension reduction and duplicate removal may improve performance for both classification and interpretability.

- Add human knowledge. Incorporating domain expertise proves helpful in improving the prediction. Exploration on new methods to better utilize existing knowledge besides using keyword matching can be another interesting research direction.

7. Conclusion

We proposed a novel multimodal machine learning approach to predict ICD-10 diagnostic codes using EMRs. We developed separate models that handle data from different modalities, including unstructured text, semi-structured text, and structured tabular data. Experiments show that our ensembled model outperforms all the baseline methods in classification. Incorporating human knowledge into machine learning models to further enhance the performance is proved effective and we would like to explore more about it in our future work. In addition, our interpretability method makes our prediction more explainable, transparent and trustworthy. The capability to extract valuable information that includes but not limited to code-specific text phrases, important lab test record and prescription, demonstrates great potential for use in real clinical practice.

References

- Combining classifiers in text categorization. In *Proceedings of the 19th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 289–297. ACM, 1996.
- Diane L Adams, Helen Norman, and Valentine J Burroughs. Addressing medical coding and billing part ii: a strategy for achieving compliance. a risk management approach for reducing coding and billing errors. *Journal of the National Medical Association*, 94(6): 430, 2002.
- Alan R Aronson. Effective mapping of biomedical text to the umls metathesaurus: the metamap program. In *Proceedings of the AMIA Symposium*, page 17. American Medical Informatics Association, 2001.
- David W Bates, Mark Ebell, Edward Gotlieb, John Zapp, and HC Mullins. A proposal for electronic medical records in us primary care. *Journal of the American Medical Informatics Association*, 10(1):1–10, 2003.
- Alex Bottle and Paul Aylin. Intelligent information: a national system for monitoring clinical performance. *Health services research*, 43(1p1):10–31, 2008.
- De Cheng, Yihong Gong, Sanping Zhou, Jinjun Wang, and Nanning Zheng. Person re-identification by multi-channel parts-based cnn with improved triplet loss function. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1335–1344, 2016.
- Thomas G Dietterich. Ensemble methods in machine learning. In *International workshop on multiple classifier systems*, pages 1–15. Springer, 2000.

- Francisco Duarte, Bruno Martins, Cátia Sousa Pinto, and Mário J Silva. Deep neural models for icd-10 coding of death certificates and autopsy reports in free-text. *Journal of biomedical informatics*, 80:64–77, 2018.
- Cyril Goutte and Eric Gaussier. A probabilistic interpretation of precision, recall and f-score, with implication for evaluation. In *European Conference on Information Retrieval*, pages 345–359. Springer, 2005.
- Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. Densely connected convolutional networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 2261–2269, 2017.
- Jin Huang and Charles X Ling. Using auc and accuracy in evaluating learning algorithms. *IEEE Transactions on knowledge and Data Engineering*, 17(3):299–310, 2005.
- Zhiheng Huang, Wei Xu, and Kai Yu. Bidirectional lstm-crf models for sequence tagging. *arXiv preprint arXiv:1508.01991*, 2015.
- Alistair EW Johnson, Tom J Pollard, Lu Shen, H Lehman Li-wei, 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.
- Ramakanth Kavuluru, Anthony Rios, and Yuan Lu. An empirical evaluation of supervised learning approaches in assigning diagnosis codes to electronic medical records. *Artificial intelligence in medicine*, 65(2):155–166, 2015.
- Yoon Kim. Convolutional neural networks for sentence classification. *CoRR*, abs/1408.5882, 2014.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Bevan Koopman, Guido Zuccon, Anthony Nguyen, Anton Bergheim, and Narelle Grayson. Automatic icd-10 classification of cancers from free-text death certificates. *International journal of medical informatics*, 84(11):956–965, 2015.
- Jeong Min Lee and Aldrian Obaja Muis. Diagnosis code prediction from electronic health records as multilabel text classification: A survey.
- Zachary C Lipton, David C Kale, Charles Elkan, and Randall Wetzell. Learning to diagnose with lstm recurrent neural networks. *arXiv preprint arXiv:1511.03677*, 2015.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013a.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119, 2013b.

- James Mullenbach, Sarah Wiegrefe, Jon Duke, Jimeng Sun, and Jacob Eisenstein. Explainable prediction of medical codes from clinical text. *arXiv preprint arXiv:1802.05695*, 2018.
- Shyamala G Nadathur. Maximising the value of hospital administrative datasets. *Australian Health Review*, 34(2):216–223, 2010.
- Suphakit Niwattanakul, Jatsada Singthongchai, Ekkachai Naenudorn, and Supachanun Wanapu. Using of jaccard coefficient for keywords similarity. In *Proceedings of the International MultiConference of Engineers and Computer Scientists*, volume 1, 2013.
- Kimberly J O’malley, Karon F Cook, Matt D Price, Kimberly Raiford Wildes, John F Hurdle, and Carol M Ashton. Measuring diagnoses: Icd code accuracy. *Health services research*, 40(5p2):1620–1639, 2005.
- G Parthiban and SK Srivatsa. Applying machine learning methods in diagnosing heart disease for diabetic patients. *International Journal of Applied Information Systems (IJ AIS)*, 3:2249–0868, 2012.
- Hude Quan, Vijaya Sundararajan, Patricia Halfon, Andrew Fong, Bernard Burnand, Jean-Christophe Luthi, L Duncan Saunders, Cynthia A Beck, Thomas E Feasby, and William A Ghali. Coding algorithms for defining comorbidities in icd-9-cm and icd-10 administrative data. *Medical care*, pages 1130–1139, 2005.
- J. R. Quinlan. Induction of decision trees. *Mach. Learn.*, 1(1):81–106, March 1986. ISSN 0885-6125. doi: 10.1023/A:1022643204877.
- Juan Ramos et al. Using tf-idf to determine word relevance in document queries. In *Proceedings of the first instructional conference on machine learning*, volume 242, pages 133–142, 2003.
- Marco Túlio Ribeiro, Sameer Singh, and Carlos Guestrin. ”why should I trust you?”: Explaining the predictions of any classifier. *CoRR*, abs/1602.04938, 2016.
- Ryan Rifkin and Aldebaro Klautau. In defense of one-vs-all classification. *J. Mach. Learn. Res.*, 5:101–141, December 2004. ISSN 1532-4435.
- Elyne Scheurwegs, Kim Luyckx, Léon Luyten, Bart Goethals, and Walter Daelemans. Assigning clinical codes with data-driven concept representation on dutch clinical free text. *Journal of biomedical informatics*, 69:118–127, 2017.
- 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.
- Benjamin Shickel, Patrick James Tighe, Azra Bihorac, and Parisa Rashidi. Deep ehr: A survey of recent advances in deep learning techniques for electronic health record (ehr) analysis. *IEEE journal of biomedical and health informatics*, 22(5):1589–1604, 2018.
- John Shore and Rodney Johnson. Axiomatic derivation of the principle of maximum entropy and the principle of minimum cross-entropy. *IEEE Transactions on information theory*, 26(1):26–37, 1980.

- Mary H Stanfill, Margaret Williams, Susan H Fenton, Robert A Jenders, and William R Hersh. A systematic literature review of automated clinical coding and classification systems. *Journal of the American Medical Informatics Association*, 17(6):646–651, 2010.
- Juan F. Escobar Stanford. Greedy function approximation : A gradient boosting machine. 2002.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. *CoRR*, abs/1512.00567, 2015.
- Grigorios Tsoumakas and Ioannis Katakis. Multi-label classification: An overview. *International Journal of Data Warehousing and Mining (IJDWM)*, 3(3):1–13, 2007.
- Pengtao Xie and Eric Xing. A neural architecture for automated icd coding. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 1066–1076, 2018.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. Character-level convolutional networks for text classification. In *Advances in neural information processing systems*, pages 649–657, 2015.

Appendix A.

Table 3: 32 ICD-10 codes and associated descriptions

ICD-10 Code	Description
I10	Essential (primary) hypertension
I50.9	Heart failure, unspecified
I48.91	Unspecified atrial fibrillation
I25.10	Atherosclerotic heart disease of native coronary artery without angina pectoris
N17.9	Acute kidney failure, unspecified
E11.9	Type 2 diabetes mellitus without complications
E78.5	Hyperlipidemia, unspecified
N39.0	Urinary tract infection, site not specified
E78.0	Pure hypercholesterolemia, unspecified
D64.9	Anemia, unspecified
E03.9	Hypothyroidism, unspecified
J18.9	Pneumonia, unspecified organism
D62	Acute posthemorrhagic anemia
R65.20	Severe sepsis without septic shock
F32.9	Major depressive disorder, single episode, unspecified
F17.200	Nicotine dependence, unspecified, uncomplicated
D69.6	Thrombocytopenia, unspecified
Z95.1	Presence of aortocoronary bypass graft
Z87.891	Personal history of nicotine dependence
I12.0	Hypertensive chronic kidney disease with stage 5 chronic kidney disease or end stage renal disease
R65.21	Severe sepsis with septic shock
Z79.4	Long term (current) use of insulin
G47.33	Obstructive sleep apnea (adult) (pediatric)
J45.909	Unspecified asthma, uncomplicated
M81.0	Age-related osteoporosis without current pathological fracture
R56.9	Unspecified convulsions
N18.6	End stage renal disease
E66.9	Obesity, unspecified
R78.81	Bacteremia
F05	Delirium due to known physiological condition
E46	Unspecified protein-calorie malnutrition
E66.01	Morbid (severe) obesity due to excess calories

Table 4: Samples of Prediction Interpretability from different modalities. Samples are displayed as in descending order and their rankings are shown in brackets. Due to space limit, only top 3 unstructured text results are included, except those mentioned in 6.2.

Interpretability Demonstration	
D64.9 Anemia, Unspecified	
Unstructured Text	‘Right Ventricular’ (1) ‘Postoperative Pleural Effusions, Anemia, Acute Renal (2) nosocomial pneumonia. Cultures eventually grew out mrsa’ (3)
LABEVENTS	‘Hematocrit’(1), ‘Hyaline Casts’(2), ‘pO2’(3), ‘Hemoglobin’(10)
PRESCRIPTION	‘Vancomycin’(1), ‘Heparin’(2), ‘Sodium Chloride 0.9% Flush’(3)
E11.9 Type 2 diabetes mellitus without complications	
Unstructured Text	‘Medical History: Copd Obesity DM II’ (1) ‘nebs changed to inhalers. Pts DM’ (2) ‘COPD, DM, Tobacco Abuse, Obesity, Iron Deficiency Anemia’ (3) ‘Metformin 850 mg po’ (5)
LABEVENTS	‘Glucose’(1), ‘Basophils’(2), ‘Lactate Dehydrogenase (LD)’(3)
PRESCRIPTION	‘Insulin’(1), ‘Metformin’(2), ‘Glyburide’(3), ‘Humulin-R Insulin’(6)
I10 Essential (primary) hypertension	
Unstructured Text	‘coumadin, HTN, COPD, Hepatocellular carcinoma’ (1) ‘diagnostic thoracentesis, urinalysis negative, urine cx’ (2) ‘subsequent enucleation, stent in pancreas’ (3)
LABEVENTS	‘Creatinine’(1), ‘Monocytes’(2), ‘Myelocytes’(3)
PRESCRIPTION	‘Atenolol’(1), ‘Sodium CHloride 0.9% Flush’(2), ‘Albuterol-Ipratropium’(3)
I50.9 Heart failure, unspecified	
Unstructured Text	‘Lisinopril’ (1), ‘Anxiety/Depression, CKD, HLD, Obesity, HTN’ (2) ‘Chronic Systolic Heart Failure’ (3)
LABEVENTS	‘Troponin T’(1), ‘Urea Nitrogen’(2), ‘Heart_rate’(3)
PRESCRIPTION	‘Furosemide’(1), ‘Carvedilol’(2), ‘Topiramate (Topamax)’(3) , ‘Lisinopril’(4)
N17.9 Acute kidney failure, unspecified	
Unstructured Text	‘Acute Kidney Injury: Multifactorial etiology’ (1), ‘Renal failure’ (2) ‘Hands, wrists, elbows, shoulders, forearm developed a’ (3) ‘IVDU, states he used heroin, barbiturates, cocaine’ (5)
LABEVENTS	‘Creatinine’(1), ‘Urea Nitrogen’(2), ‘Monocytes’(3)
PRESCRIPTION	‘NaHCO3’(1), ‘Iso-Osmotic Dextrose’(2), ‘Phytonadione’(3)