

# Multi-Label Learning from Medical Plain Text with Convolutional Residual Models

**Xinyuan Zhang**

*Department of Electrical and Computer Engineering  
Duke University  
Durham, NC, USA*

XY.ZHANG@DUKE.EDU

**Ricardo Henaó**

*Department of Biostatistics and Bioinformatics  
Duke University  
Durham, NC, USA*

RICARDO.HENAO@DUKE.EDU

**Zhe Gan**

*Department of Electrical and Computer Engineering  
Duke University  
Durham, NC, USA*

ZHE.GAN@DUKE.EDU

**Yitong Li**

*Department of Electrical and Computer Engineering  
Duke University  
Durham, NC, USA*

YITONG.LI@DUKE.EDU

**Lawrence Carin**

*Department of Electrical and Computer Engineering  
Duke University  
Durham, NC, USA*

LCARIN@DUKE.EDU

## Abstract

Predicting diagnoses from Electronic Health Records (EHRs) is an important medical application of multi-label learning. We propose a convolutional residual model for multi-label classification from *doctor notes* in EHR data. A given patient may have multiple diagnoses, and therefore multi-label learning is required. We employ a Convolutional Neural Network (CNN) to encode plain text into a fixed-length sentence embedding vector. Since diagnoses are typically correlated, a deep residual network is employed on top of the CNN encoder, to capture label (diagnosis) dependencies and incorporate information directly from the encoded sentence vector. A real EHR dataset is considered, and we compare the proposed model with several well-known baselines, to predict diagnoses based on doctor notes. Experimental results demonstrate the superiority of the proposed convolutional residual model.

## 1. Introduction

Machine learning is playing an increasingly important role in medical applications. The widespread use of Electronic Health Records (EHRs) has made it possible to collect massive amounts of data from patients and health providers. Machine learning technology can help clinical researchers (*i*) discover hidden patterns from massive EHR data, and (*ii*) develop predictive models to assist with clinical decision making. One important application of multi-label learning in the medical domain is to predict diagnoses given features from the EHR.

Multi-label learning is a supervised classification framework in which multiple target labels can be assigned to one instance. For example, in our motivating doctor notes dataset, one patient can

be associated with multiple diagnoses simultaneously (comorbidities), *e.g.*, “fever”, “cough”, and “viral infection”, as one patient may suffer from several related illnesses. One challenge of multi-label learning is modeling label dependencies by realizing that labels are often correlated. In the above example, and assuming that the {fever, cough, viral infection} combination is likely, if a patient manifests fever and cough, then the probability of also having a viral infection is likely to increase accordingly.

Traditional multi-label learning methods such as *one-versus-all* and *one-versus-one* Zhang and Zhou (2014) assume that labels are independent of each other. Recent work in multi-label learning focuses on exploiting label correlation to improve classification performance. A natural approach consists of using embedding-based models to project label vectors onto a low dimensional space while capturing dependencies, thus reducing the “effective” number of labels Bhatia et al. (2015). However, in practice, embedding-based methods can result in subpar performance due to the loss of information during the embedding procedure. Tree-based approaches achieve faster prediction by recursively partitioning labels into tree-structured groups Prabhu and Varma (2014). However, errors made at the upper levels of the hierarchy cannot be corrected at bottom levels, which often leads to a loss in overall predictive accuracy. In addition, most models referenced above use *bag-of-words* sentence representations. The main shortcoming of bag-of-words models is their inability to capture local context (semantics). For example, “buy used cars” and “purchase old automobiles” are represented by orthogonal (unrelated) vectors in a bag-of-words representation, but in fact they are semantically identical. The order of words is not respected in bag-of-words representations. For example, “he is healthy” and “is he healthy” have exactly the same bag-of-words representation but have vastly different meanings. Long Short-Term Memory (LSTM) Hochreiter and Schmidhuber (1997) is a widely used approach to estimate fixed-length vectorial representations to capture sentence meaning. However, an LSTM typically becomes ineffective when modeling very long sentences or paragraphs, which unfortunately are typical in doctor notes.

In this paper we consider a real EHR-based dataset, where the goal is to predict diagnoses based on plain-text doctor notes. This scenario is not a typical sentence classification task, in the sense that each note is composed of several sentences (a paragraph). The average note length is  $101 \pm 48$  words. To evaluate the proposed model, we compare it with a number of related methods on our doctor notes dataset. Experiments show that our convolutional residual model outperforms all other competing methods. The superiority of the proposed model indicates that: (i) compared to bag-of-words models, the CNN is more effective at sentence encoding, by leveraging semantics and word ordering, and (ii) deep residual networks successfully capture label dependencies, thus delivering significantly improved multi-label classification performance.

**Technical Significance** We develop a convolutional residual model to address this multi-label plain text learning problem. CNN models have been increasingly used for Natural Language Processing (NLP) applications, and have achieved excellent results on both supervised Kalchbrenner et al. (2014); Hu et al. (2014) and unsupervised Gan et al. (2017) tasks. Here, we employ a CNN as the sentence (doctor’s note) encoder, due to its excellent performance at identifying sentence structure, especially in long noisy sentences. A deep residual network He et al. (2016) is added on top of the CNN encoder, to capture label correlations and incorporate information from the encoded sentence vector (CNN’s output). This is achieved using shortcut connections between layers Bishop (1995); Venables and Ripley (2013). A common problem in deep networks is that performance

saturates, then degrades rapidly as the network grows deeper. Deep residual networks mitigate this problem by having stacked layers that fit a residual mapping.

**Clinical Relevance** This work focuses on predicting diagnoses given a plain doctor’s note on a patient’s presentation, symptoms, medical history, etc. However, leveraging the dependencies of multiple comorbidities in an heterogeneous population of patients remains a very challenging problem. Exploiting diagnoses correlation (co-occurrence) is essential when approaching this predictive problem with multi-label classification methods. Further, a large number of typos, medical jargon, and non-standard abbreviations make the notes considerably noisy and heterogeneous. Our goal is to build a classifier as robust as possible while minimizing the preprocessing burden, to address the above problems and simplify the implementation of the predictive model in our local Medical System.

## 2. Related Work

There is an extensive literature on applying NLP techniques to medical-domain tasks. Notably, Jagannatha and Yu (2016) applied RNN-based sequence labeling in phrase detection of medical text. Kuo et al. (2016) built an NLP ensemble pipeline to integrate two systems cTAKES Savova et al. (2010) and MetaMap Aronson and Lang (2010) for biomedical data-element extraction from clinical text. Recently, Li et al. (2017) employed CNN on textual admission information to predict medical diagnosis. However, none of these approaches consider the (practical) problem of diagnosing patients affected by multiple illnesses, which belongs to the larger class of multi-label learning methods.

Existing multi-label classification methods are commonly grouped into two broad categories. One is problem-transformation methods, that convert a multi-label learning problem into several binary classification tasks by using existing techniques such as Binary Relevance (BR) Tsoumakas et al. (2009), Label Power-set (LP) Tsoumakas and Katakis (2006) and the pair-wise method Wu et al. (2004). The other category consists of algorithm-adaptation methods, which extend specific supervised algorithms to deal with multi-label data. For instance, Li et al. (2015) leveraged the conditional Restricted Boltzmann Machine (RBM) to multi-label tasks. Our model belongs to the latter group, in the sense that we extend CNN-based classifiers.

Recently, multi-label learning has been employed for predicting diagnoses based on clinical data. Zhao et al. (2013) developed the Ensemble of Sampled Classifier Chains (ESCC) algorithm by exploiting disease label relationships, to classify clinical text according to their specified disease labels. Li et al. (2016) performed multi-label classification for health and disease risk prediction using a Multi-Label Problem Transformation Joint Classification (MLPTJC) method. However, none of these methods use plain text data nor CNN-based representations.

## 3. Methods

We use CNNs as a fixed-length feature extractor from word sequences, *i.e.*, doctor notes in our motivating dataset. Doctor notes are in general composed of several sentences. Here we treat each note as a single “meta” sentence, with end-of-sentence tokens located at the end of each actual sentence. Three different classifiers based on the CNN-encoded sentence features are described and evaluated, including our convolutional residual model described in Section 3.4.

A sentence  $S$  with  $T$  words  $\{w_1, w_2, \dots, w_T\}$  is mapped into a set of  $k$ -dimensional real-valued vectors  $\{e_1, e_2, \dots, e_T\}$  using a word embedding. Word vectors are initialized using *word2vec* which is pre-trained on 100 billion words from Google News, including numbers and special characters, using a continuous bag-of-words architecture Mikolov et al. (2013). This embedding matrix is further refined using medical text during training.

### 3.1 CNN Sentence Encoder

Based on the word embedding, sentence  $S$  is represented as a  $\mathbf{X} \in \mathbb{R}^{k \times T}$  matrix, assembled by concatenating word vectors  $e_1, e_2, \dots, e_T$ , i.e.,

$$\mathbf{X} = e_1 \oplus e_2 \oplus \dots \oplus e_T,$$

where the  $i$ -th column of  $\mathbf{X}$  is the embedding vector corresponding to word  $w_i$ .

The CNN from Kim (2014); Collobert et al. (2011) with  $\mathbf{X}$  as input is utilized as the sentence encoder. Given a filter  $\Lambda \in \mathbb{R}^{k \times t}$  with a window of  $t$  words, we produce feature map  $\mathbf{g} = [g_1 \ g_2 \ \dots \ g_{T-t+1}]$  by

$$\mathbf{g} = f(\mathbf{X} * \Lambda + \epsilon), \quad (1)$$

where  $\mathbf{g} \in \mathbb{R}^{T-t+1}$ ,  $f(\cdot)$  is a nonlinear function such as the hyperbolic tangent used in this paper,  $*$  is the convolution operator and  $\epsilon \in \mathbb{R}^{T-t+1}$  is a bias term.

Note the obtained feature map,  $\mathbf{g}$ , of length  $T - t + 1$  depends of the sentence length. To deal with this issue, the *max*-pooling operator is applied on the feature map. By taking the maximum value  $\hat{g} = \max\{\mathbf{g}\}$ , only the most salient sentence feature corresponding to the filter,  $\Lambda$ , is captured for sentence  $S$  via  $\mathbf{X}$ .

Equation (1) describes how the model uses one filter to extract one sentence feature. In practice, multiple filters of varying window sizes act as linguistic feature detectors, whose goal is to recognize specific classes of  $n$ -grams. For instance, if we set  $v$  filters, where each filter has  $w$  variant window sizes, then the resulting encoded sentence representations is a  $vw$ -dimensional vector.

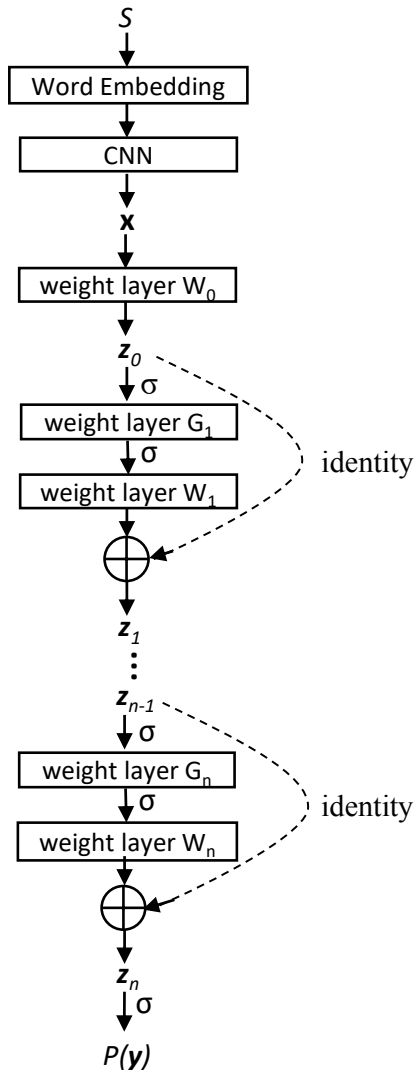


Figure 1: The building blocks of an  $n$ -layer convolutional residual model.

### 3.2 CNN Classifiers

Let  $\mathbf{x} \in \mathbb{R}^{vw}$  be the encoded sentence vector. In multi-label learning, the label layer  $\mathbf{y} = [y^{(1)} \dots y^{(L)}] \in \{0, 1\}^L$  represents the underlying true label vector (diagnoses in our case), where  $L$  is the total number of labels and  $y_i = 1$  indicates the existence of  $i$ -th label. Given a sentence vector  $\mathbf{x}$ , the probability of all labels  $P(\mathbf{y}) = [P(y^{(1)} = 1) \dots P(y^{(L)} = 1)]$  can be expressed as

$$P(\mathbf{y}) = \sigma(\mathbf{z}_0), \quad \mathbf{z}_0 = \mathbf{W}_0 \mathbf{x} + \mathbf{b}_0, \quad (2)$$

where  $P(\mathbf{y}) \in [0, 1]^L$ ,  $\sigma(\cdot)$  denotes the sigmoid link function,  $\sigma(x) = 1/(1 + \exp(-x))$ ,  $\mathbf{W}_0 \in \mathbb{R}^{L \times vw}$  is the (classification) weights matrix, and  $\mathbf{b}_0 \in \mathbb{R}^L$  is the bias term.

### 3.3 Conditional RBM Models

The CNN classifier in (2) treats each label independently, as  $L$  binary classification problems. However, in real multi-label tasks the labels are usually correlated with each other. The model proposed by Li et al. (2015) for bag-of-words representations uses a Restricted Boltzmann Machine (RBM) to capture this high-order label dependency.

A latent layer  $\mathbf{h} \in \{0, 1\}^J$  with  $J$  hidden units is added above the label layer  $\mathbf{y} \in \{0, 1\}^L$ . Conditioning on the input feature layer  $\mathbf{x}$ , the encoded sentence vector in our case, layers  $\mathbf{y}$  and  $\mathbf{h}$  form a standard restricted Boltzmann machine. The Conditional RBM (CRBM) model is specified via parameters:  $\mathbf{W} \in \mathbb{R}^{L \times vw}$ ,  $\mathbf{G} \in \mathbb{R}^{L \times J}$ ,  $\mathbf{b} \in \mathbb{R}^{L \times 1}$ ,  $\mathbf{c} \in \mathbb{R}^{J \times 1}$ . The conditional marginal likelihood is defined as

$$P(\mathbf{y}|\mathbf{x}) = \frac{\sum_{\mathbf{h}} \exp[-E_{\text{con}}(\mathbf{y}, \mathbf{x}) - E_{\text{rbm}}(\mathbf{y}, \mathbf{h})]}{Z(\mathbf{x})},$$

where  $Z(\mathbf{x})$  is the normalization factor, and  $E_{\text{con}}(\mathbf{y}, \mathbf{x})$  and  $E_{\text{rbm}}(\mathbf{y}, \mathbf{x})$  are two energy functions such that

$$\begin{aligned} Z(\mathbf{x}) &= \sum_{\mathbf{y}} \sum_{\mathbf{h}} \exp[-E_{\text{con}}(\mathbf{y}, \mathbf{x}) - E_{\text{rbm}}(\mathbf{y}, \mathbf{h})], \\ E_{\text{con}}(\mathbf{y}, \mathbf{x}) &= -\mathbf{y}^\top \mathbf{W} \mathbf{x}, \\ E_{\text{rbm}}(\mathbf{y}, \mathbf{h}) &= -\mathbf{y}^\top \mathbf{G} \mathbf{h} - \mathbf{y}^\top \mathbf{b} - \mathbf{c}^\top \mathbf{h}. \end{aligned}$$

Given training data, the CRBM model is optimized by maximizing the conditional marginal likelihood,  $P(\mathbf{y}|\mathbf{x})$ .

According to the conditional independence structure of RBM models, the local conditional probabilities can be computed as

$$\begin{aligned} P(h^{(j)} = 1|\mathbf{y}, \mathbf{x}) &= \sigma(\mathbf{y}^\top \mathbf{G}_{:j} + \mathbf{c}_j), \\ P(y^{(l)} = 1|\mathbf{h}, \mathbf{x}) &= \sigma(\mathbf{G}_l \mathbf{h} + \mathbf{b}_l + \mathbf{W}_l \mathbf{x}). \end{aligned} \quad (3)$$

where  $\mathbf{G}_{:j}$  and  $\mathbf{G}_l$  are column  $j$  and row  $l$  of  $\mathbf{G}$ , respectively.

### 3.4 Convolutional Residual Models

Motivated by the local conditional probabilities of CRBM models in (3), we add feedforward neural networks with shortcut connections on top of the CNN encoder in Section 3.1. Shortcut connections

---

**Algorithm 1:** Convolutional Residual Model.

---

**Input** : Sentence  $S$  with  $T$  words.  
**Output:**  $P(\mathbf{y})$ : Probability of each label.  
 Encode  $S$  into vector  $\mathbf{x} \in \mathbb{R}^{vw}$  using CNN.  
 Initialize base layer:  $\mathbf{z}_0 = \mathbf{W}_0\mathbf{x} + \mathbf{b}_0$ .  
**for**  $i = 1, 2, \dots, n$  **do**  
      $\mathbf{q}_i = \sigma(\mathbf{z}_{i-1})\mathbf{G}_i + \mathbf{c}_i$ .  
      $\mathbf{z}_i = \mathbf{W}_0\mathbf{x} + \mathbf{b}_i + \sum_{t=1}^i \mathbf{W}_t\sigma(\mathbf{q}_t)$ .  
**end**  
 $P(\mathbf{y}) = \sigma(\mathbf{z}_n)$ .

---

are defined as those skipping one or more layers. The idea is to incorporate information from both the sentence layer via the encoded sentence vector,  $\mathbf{x}$ , and the label layer,  $\mathbf{y}$ , the latter to capture label dependencies. These two components form essentially a *deep residual network*. Shortcut connections capture the predictive interactions between the encoded sentence vectors and the output labels, while the stacked feedforward neural network captures the correlations in the label layer. Shortcut connections enable the stacked layers to fit a residual mapping, which avoids model degradation as the network depth (number of layers) increases He et al. (2016). In this paper, shortcut connections are constructed as identity mappings, thus the encoded sentence vectors are directly added to the outputs of the stacked layers.

Our complete convolutional residual model is composed of the CNN encoder in Section 3.1 and the residual classifier described above. The building blocks of an  $n$ -layer residual classifier are shown in Figure 1, in which  $\oplus$  denotes element-wise addition.  $\mathbf{x}$  represents the CNN encoded sentence vector and  $\mathbf{W}_0$  is the weight in Equation 2. We set sigmoid link functions,  $\sigma(\cdot)$ , as nonlinearities on every layer, and the biases are omitted for simplicity. Further,  $n$  is the number of residual layers and  $h_1, \dots, h_i, \dots, h_n$  are the number of hidden units of each hidden layer.  $\mathbf{W}_0 \in \mathbb{R}^{L \times vw}$ ,  $\mathbf{W}_1, \dots, \mathbf{W}_n \in \mathbb{R}^{L \times h_i}$ ,  $\mathbf{b}_0, \mathbf{b}_1, \dots, \mathbf{b}_n \in \mathbb{R}^L$ ,  $\mathbf{G}_1, \dots, \mathbf{G}_n \in \mathbb{R}^{L \times h_i}$ , and  $\mathbf{c}_1, \dots, \mathbf{c}_n \in \mathbb{R}^{h_i}$  are parameters to be learned, namely, weight matrices and bias vectors for layers  $0, \dots, n$ . Identity shortcut connections introduce neither additional parameters nor computational complexity. This allows us to fairly compare plain stacked classifiers and residual classifiers. The algorithm for our *Convolutional Residual Model* for multi-label tasks is presented as Algorithm 1. A plain stacked classifier (without shortcut connections) correspond to the architecture in Figure 1 without the identity connections (dashed lines).

The residual and plain stacked classifiers have exactly the same number of parameters. Let  $\Theta$  be the set of parameters of both the CNN encoder and the residual classifier. We wish to find the optimal set of parameters that minimize the cross-entropy loss function, expressed as

$$\arg \min_{\Theta} \frac{1}{L} \sum_{i=1}^L y^{(i)} \log P(y^{(i)} = 1) + (1 - y^{(i)}) \log[1 - P(y^{(i)} = 1)].$$

The parameters,  $\Theta$ , of the entire network are jointly optimized by stochastic gradient descent with back-propagation.

Dataset	$L$	Label Type	Controls	$N$	$ V $	$l$	$N_V$
<i>ehr25</i>	25	Single-Label	No	7897	19596	410	717
<i>ehr1000</i>	1000	Multi-Label	No	44473	54215	600	4447
<i>ehr64-all</i>	64	Multi-Label	Yes	50128	58359	600	5012

Table 1: Dataset summary.  $L$ : the number of labels.  $l$ : maximum sentence length.  $N$ : dataset size.  $|V|$ : vocabulary size.  $N_V$ : validation set size.

#### 4. Doctor Notes Data

This work is motivated by a real EHR cohort collected from patients at Duke Hospital. The input plain text is a doctor’s note on a patient’s history of present illness, which briefly describes in the doctor’s words information about the patient with regard to presentation, symptoms, medical history, *etc.* The outputs (labels) are free-text discharge diagnoses, medications, and dispositions. In this paper we focus on diagnoses as the target of interest. Note that each patient could have multiple diagnoses simultaneously. Furthermore, we do not preprocess the plain text data in any way, which means we retain typos, medical jargon and non-standard abbreviations in the notes.

It is important to consider that discharge diagnoses adjudication does not occur immediately after the doctor’s note has been written. In fact, in our data, about 50% of the patients have to wait for at least 7 hours to be diagnosed, 25% have to wait for at least 18 hours and only 12% of patients have to wait less than 2 hours. This time is usually spent on laboratory tests and medical procedures required to confirm the discharge diagnoses. Hence, our predictions are of *future* diagnoses based on doctor notes, which presents doctors with the most likely diagnoses to guide orders and hopefully improve care.

Below we present an example of a doctor’s note from our dataset. Note the use of special characters, abbreviations (hx, s/p, po, pt, *etc.*) and typos.

**Note:** 76 yo woman with hx of cad s/p stent, degen disc disease, ? dementia, ? past tia, htn, and anxiety d/o seen multiple times in the ed for lapses in personal hygeine and “spells” now brought in by her daughter after “hollering and screaming” this morning at home (where she lives with another daughter). pt denies anger or being upset this morning, has no current complaints. no difficulty with po intake or taking medications, per pt. continues to endorse pain in l lower abdomen yesterday, which is now resolved. pt states that she has had diarrhea x 1 week or so, but no difficulty or pain with urination. no fevers, chills, or other symptoms. no prior hx of similar problem.

**Diagnoses:** anxiety, hypertension.

It is well understood that diagnoses exhibit a rich correlation (co-occurrence) structure, *e.g.*, anxiety and hypertension as in the example above. Leveraging such correlations may be very beneficial in practice, when attempting to make multiple predictions for a single patient based solely on doctor’s notes.

The dataset has over 50,128 doctor notes and 8,279 free-text diagnoses. However, 7279 out of 8279 diagnoses have less than 10 instances. Since a considerable amount of diagnoses do not have enough number of samples to build and reliably evaluate a supervised model, we generate three sub-datasets by focusing on the most common diagnoses.

Table 1 shows a summary of the datasets being considered. Different datasets are used for different experiments. *ehr25* consists of patients with a single diagnosis, while restricting to the most common 25 diagnoses (each diagnosis has at least 140 instances). *ehr1000* consists of patients with at least one diagnosis, while restricting to the most common 1000 diagnoses (each diagnosis has at least 10 instances). *ehr64-all* consists of all doctors’ notes including patients with the most common 64 diagnoses (26071 instances) and “control patients” with no labels (24057 instances). Note that the “control patients” group has diagnoses information available but they are not within the top 64 set.

## 5. Experiments

Our experiments are conducted on different subsets of the cohort described in Section 4. We compare our convolutional residual models with  $n = \{1, 2, 4, 8\}$  residual layers against five baselines on the *ehr1000* dataset. Then we use the trained parameters of the CNN encoder on the single-labeled *ehr25* dataset to visualize CNN encoded doctor’s notes and their corresponding diagnosis using *t*-SNE Maaten and Hinton (2008). Finally, we compare the results between a plain CNN classifier and convolutional residual model with 8 layers on the *ehr64-all* dataset. Our model aims to encode plain text to sentence vectors using CNNs. However, to the best of the author’s knowledge, all of publicly available multi-label text classification datasets have been preprocessed as numerical vectors such as bag-of-words. Besides, generating artificial data for this task is extremely difficult due to the challenging nature of generating synthetic natural language. Therefore, we focus on the real EHR dataset to find the best approach for the intended application task.

For the word embedding, we specify an embedding matrix  $W_{\text{emb}} \in \mathbb{R}^{k \times |V|}$  for each word in sentence  $S$ , where  $V$  is a fixed-sized vocabulary of size  $|V|$ , and each word is represented by a  $k$ -dimensional real-valued word vector. In *word2vec*, the dimensionality  $k$  is set to 300. For words that do not occur in the *word2vec* embedding matrix, we *learn* a corresponding 300-dimensional word vector (initialized with entries randomly sampled from  $[-0.25, 0.25]$ , such that the new vectors have approximately the same variance as the pre-trained vectors). Hence, embeddings are learned for common misspelled words and typos, typically found in doctor notes. The subset of words in the validation set that do not appear in the training set are replaced by the “unknown” token from *word2vec*, because they only appear once in the entire dataset. This subset accounts for 0.5% per-sentence words or up to one word per validation sentence. For the CNN encoder, we set filter windows of sizes  $h = \{3, 4, 5\}$  with 100 filters each, thus each sentence is encoded as a 300-dimensional vector.

All models are implemented in Theano Bastien et al. (2012), using a NVIDIA Titan X GPU with 12 GB memory. The word embedding matrix is initialized by *word2vec*. All parameters of the model, namely weights of the CNN encoder and weights in the residual multi-layer classifier are initialized from a uniform distribution with support  $[-0.01, 0.01]$ . All bias terms are initialized to zero. We use Adam Kingma and Ba (2014) with learning rate 0.0002 for the optimization procedure. 10% of the dataset is used as validation set. Early stopping is employed to avoid overfitting. The minibatch size is set to 50. We use dropout Srivastava et al. (2014) with a probability of 0.50 for the CNN encoder. We will make the source codes publicly available.



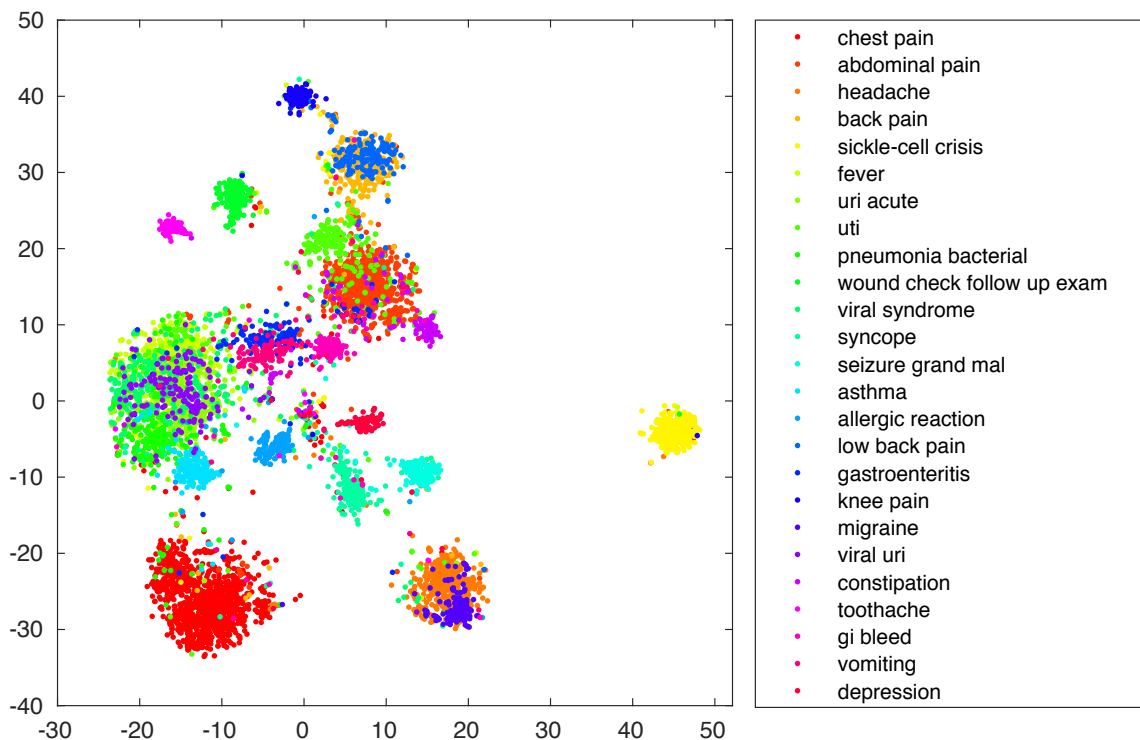


Figure 2: Visualization of CNN-encoded doctor's notes on *ehr25* using *t*-SNE.

## 5.1 Baselines

**SLEEC** is an embedding-based multi-label learning technique proposed by Bhatia et al. (2015). This bag-of-words algorithm can be used for large-scale problems by learning embeddings which preserve pairwise distances between nearest label vectors.

**FastXML** is a tree-based multi-label learning technique proposed by Prabhu and Varma (2014). This bag-of-words algorithm can make fast predictions by directly optimizing an nDCG-based ranking loss function.

**Bi-LSTM** A Bidirectional LSTM Graves et al. (2013) consisting of paired LSTMs connected in opposite directions is used for sentence encoding. Then we build the logistic regression model in (2) on top of it as classifier.

**CNN** A logistic regression model in (2) on top of the CNN-encoded sentence features as described in Section 3.2.

**CRBM** A conditional restricted Boltzmann machine model as in Section 3.3, using CNN-encoded sentence vectors as input features. The parameters of the CNN encoder are trained by a CNN classifier.

**Convolutional Plain Models** An 8-layer convolutional plain model (no shortcut connections). The model has exactly the same set of parameters as an 8-layer convolutional residual model.

Model	P@1(%)	P@3(%)	P@5(%)	N@3(%)	N@5(%)	AUC(%)
SLEEC	36.36	19.74	14.03	43.00	46.64	—
fast-XML	57.64	29.38	20.04	65.66	69.52	91.84
Bi-LSTM	55.49	28.35	19.67	73.53	73.52	92.48
CNN	58.87	30.59	20.77	78.91	78.88	93.38
CRBM	54.81	26.69	18.31	60.93	64.62	92.59
CNN-Plain-8layer	8.99	8.19	7.19	20.68	20.67	50.04
CNN-Res-1layer	59.98	31.04	21.11	79.71	79.67	94.56
CNN-Res-2layer	60.08	31.04	21.12	79.77	79.73	94.61
CNN-Res-4layer	60.28	31.20	<b>21.32</b>	80.12	80.09	94.78
CNN-Res-8layer	<b>60.30</b>	<b>31.21</b>	21.27	<b>80.21</b>	<b>80.17</b>	<b>94.89</b>

Table 2: Quantitative results using several metrics for doctor-notes multi-label classification task on *ehr1000*. Scores correspond to validation data. The best results are in bold. For SLEEC we do not report an AUC value because the method does not return a classification score.

## 5.2 Evaluation Metrics

**Precision@ $k$**  Precision at  $k$  is a popular evaluation metric for multi-label classification problems. Given the ground truth label vector  $\mathbf{y} \in \{0, 1\}^L$  and the prediction  $\hat{\mathbf{y}} \in [0, 1]^L$ ,  $P@k$  is defined as

$$P@k := \frac{1}{k} \sum_{l \in \text{rank}_k(\hat{\mathbf{y}})} y^{(l)}.$$

Precision at  $k$  performs a sentence-wise evaluation that counts the fraction of correct predictions in the top  $k$  scoring labels.

**nDCG@ $k$**  normalized Discounted Cumulative Gain (nDCG) at rank  $k$  is a family of ranking measures widely used in multi-label learning. DCG is the total gain accumulated at a particular rank  $p$ , which is defined as

$$DCG@k := \sum_{l \in \text{rank}_k(\hat{\mathbf{y}})} \frac{y^{(l)}}{\log(l+1)}.$$

Then normalizing DCG by the value at rank  $k$  of the ideal ranking gives

$$N@k := \frac{DCG@k}{\sum_{l=1}^{\min(k, \|\mathbf{y}\|_0)} \frac{1}{\log(l+1)}}$$

Here, nDCG at  $k$  ( $N@k$ ) performs a sentence-wise evaluation.

**AUC** The area under the receiving operating characteristic curve (AUC) is the probability that a classifier will rank a randomly chosen positive instance higher than a randomly chosen negative one. Though originally defined for binary problems, the labels can be represented as a  $N \times L$  binary matrix. We estimate the AUCs of each label individually, then take the average. AUC is then a label-wise evaluation metric averaged over sentences.

## 5.3 Results

The precision at 1, 3, 5, nDCG at 3, 5, and AUC scores for all methods on the *ehr1000* dataset are shown in Table 2. Our models show the best multi-label classification performance compared to all other competing methods. The model benefits from adding residual layers in particular when it is shallow. As can be seen, convolutional residual models with 4 or 8 residual layers have better quantitative results than those with 1 or 2 layers, while the performance difference between models with 4 layers and 8 layers is not significant. Moreover, the model converges faster as the number of residual layers increases, e.g., it takes 38 and 25 epochs for CNN-Res-1layer and CNN-Res-8layer to converge, respectively. Note that the 8-layer convolutional plain model, which has the same number of parameters as the 8-layer convolutional residual model, fails completely, indicating that the performance improvement of the proposed model is not simply due to the number of parameters increasing with the number of layers. Models using CNN-encoded sentence vectors as input features generally perform better than those using bag-of-words representations and LSTM-encoded sentence vectors as input features. This demonstrates that the CNN encoder is better at capturing sentence structures for the long and noisy doctors’ notes in our data.

We use the trained CNN parameters to encode single-label sentences in *ehr25*. In Figure 2, the CNN-encoded sentence vectors (300-dimensional) are projected into a 2-dimensional space using t-SNE. As can be seen from the visualization plot, the encoded doctors’ notes with the same diagnoses are generally clustered together. Even for those regions containing a mix of colors, they usually have similar diagnoses, such as headache and migraine, back pain and low back pain, *etc.*

Finally, we compare the convolutional residual model with 8 layers against a CNN classifier on the *ehr64-all* dataset, in which about half of doctors’ notes are from “control patients”. As shown in Figure 3, our proposed model is not only more accurate on multi-label prediction, but some diagnoses exhibit significant performance gains. The AUC scores of 59 (of 64) diagnoses (red bars) are improved in the 8-layer residual network compared to the CNN classifier. For the remaining 5 (of 64) diagnoses (dark bars) CNN is only marginally better than the convolutional residual model. The top 5 diagnoses that increase the most in Figure 3 are hyperglycemia, hypotension, hyperkalemia, uti, and fever, which are also most likely to have comorbidities in our dataset. These results demonstrate that our residual model can improve prediction performance by leveraging diagnosis dependencies.

## 6. Conclusion

We developed a novel convolutional residual model for multi-label text learning. A CNN performs convolution and pooling operations to encode the input sentence into a fixed-length feature vector. Motivated by the local conditional probabilities in CRBM models, we proposed a feedforward neu-

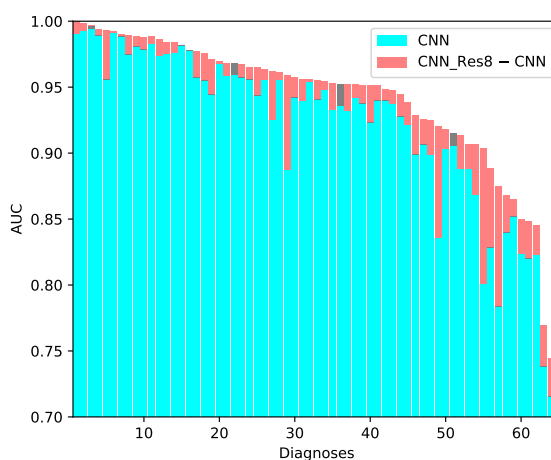


Figure 3: AUC Comparison between a CNN classifier and a 8-layer convolutional residual model on *ehr64-all*. The mean AUC of the CNN classifier and the 8-layer convolutional residual model are  $0.923 \pm 0.063$  and  $0.942 \pm 0.051$ , respectively.

ral network with shortcut connections on top of the CNN encoder as classifier. This classification structure forms a deep residual network to combine information from the encoded sentence vector and label dependencies captured from the label layer. We presented experiments on a new dataset to predict diagnoses based on plain text from doctor notes. Experimental results demonstrated the superiority of the proposed model for text embedding and multi-label learning.

## References

- Alan R Aronson and François-Michel Lang. An overview of metamap: historical perspective and recent advances. *Journal of the American Medical Informatics Association*, 17(3):229–236, 2010.
- Frédéric Bastien, Pascal Lamblin, Razvan Pascanu, James Bergstra, Ian Goodfellow, Arnaud Bergeron, Nicolas Bouchard, David Warde-Farley, and Yoshua Bengio. Theano: new features and speed improvements. *arXiv preprint arXiv:1211.5590*, 2012.
- Kush Bhatia, Himanshu Jain, Purushottam Kar, Manik Varma, and Prateek Jain. Sparse local embeddings for extreme multi-label classification. In *Advances in Neural Information Processing Systems*, pages 730–738, 2015.
- Christopher M Bishop. *Neural networks for pattern recognition*. Oxford university press, 1995.
- Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. Natural language processing (almost) from scratch. *Journal of Machine Learning Research*, 12(Aug):2493–2537, 2011.
- Zhe Gan, Yunchen Pu, Ricardo Henao, Chunyuan Li, Xiaodong He, and Lawrence Carin. Learning generic sentence representations using convolutional neural networks. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2390–2400, 2017.
- Alex Graves, Navdeep Jaitly, and Abdel-rahman Mohamed. Hybrid speech recognition with deep bidirectional lstm. In *Automatic Speech Recognition and Understanding (ASRU), 2013 IEEE Workshop on*, pages 273–278. IEEE, 2013.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- Baotian Hu, Zhengdong Lu, Hang Li, and Qingcai Chen. Convolutional neural network architectures for matching natural language sentences. In *Advances in neural information processing systems*, pages 2042–2050, 2014.
- Abhyuday N Jagannatha and Hong Yu. Structured prediction models for rnn based sequence labeling in clinical text. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing. Conference on Empirical Methods in Natural Language Processing*, volume 2016, page 856. NIH Public Access, 2016.

- Nal Kalchbrenner, Edward Grefenstette, and Phil Blunsom. A convolutional neural network for modelling sentences. *arXiv preprint arXiv:1404.2188*, 2014.
- Yoon Kim. Convolutional neural networks for sentence classification. *arXiv preprint arXiv:1408.5882*, 2014.
- Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Tsung-Ting Kuo, Pallavi Rao, Cleo Maehara, Son Doan, Juan D Chaparro, Michele E Day, Claudiu Farcas, Lucila Ohno-Machado, and Chun-Nan Hsu. Ensembles of nlp tools for data element extraction from clinical notes. In *AMIA Annual Symposium Proceedings*, volume 2016, page 1880. American Medical Informatics Association, 2016.
- Christy Li, Dimitris Konomis, Graham Neubig, Pengtao Xie, Carol Cheng, and Eric Xing. Convolutional neural networks for medical diagnosis from admission notes. *arXiv preprint arXiv:1712.02768*, 2017.
- Runzhi Li, Hongling Zhao, Yusong Lin, Andrew Maxwell, and Chaoyang Zhang. Multi-label classification for intelligent health risk prediction. In *Bioinformatics and Biomedicine (BIBM), 2016 IEEE International Conference on*, pages 986–993. IEEE, 2016.
- Xin Li, Feipeng Zhao, and Yuhong Guo. Conditional restricted boltzmann machines for multi-label learning with incomplete labels. In *AISTATS*, pages 635–643, 2015.
- Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(Nov):2579–2605, 2008.
- 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, 2013.
- Yashoteja Prabhu and Manik Varma. Fastxml: A fast, accurate and stable tree-classifier for extreme multi-label learning. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 263–272. ACM, 2014.
- Guergana K Savova, James J Masanz, Philip V Ogren, Jiaping Zheng, Sunghwan Sohn, Karin C Kipper-Schuler, and Christopher G Chute. Mayo clinical text analysis and knowledge extraction system (ctakes): architecture, component evaluation and applications. *Journal of the American Medical Informatics Association*, 17(5):507–513, 2010.
- Nitish Srivastava, Geoffrey E Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
- Grigorios Tsoumakas and Ioannis Katakis. Multi-label classification: An overview. *International Journal of Data Warehousing and Mining*, 3(3), 2006.
- Grigorios Tsoumakas, Ioannis Katakis, and Ioannis Vlahavas. Mining multi-label data. In *Data mining and knowledge discovery handbook*, pages 667–685. Springer, 2009.

William N Venables and Brian D Ripley. *Modern applied statistics with S-PLUS*. Springer Science & Business Media, 2013.

Ting-Fan Wu, Chih-Jen Lin, and Ruby C Weng. Probability estimates for multi-class classification by pairwise coupling. *Journal of Machine Learning Research*, 5(Aug):975–1005, 2004.

Min-Ling Zhang and Zhi-Hua Zhou. A review on multi-label learning algorithms. *IEEE transactions on knowledge and data engineering*, 26(8):1819–1837, 2014.

Rui-Wei Zhao, Guo-Zheng Li, Jia-Ming Liu, and Xiao Wang. Clinical multi-label free text classification by exploiting disease label relation. In *Bioinformatics and Biomedicine (BIBM), 2013 IEEE International Conference on*, pages 311–315. IEEE, 2013.