



# HHS Public Access

Author manuscript

ACM BCB. Author manuscript; available in PMC 2017 July 21.

Published in final edited form as:

ACM BCB. 2015 September ; 2015: 258–267. doi:10.1145/2808719.2808746.

## Convolutional Neural Networks for Biomedical Text Classification: Application in Indexing Biomedical Articles

**Anthony Rios** and

Department of Computer Science, University of Kentucky, Lexington, Kentucky

**Ramakanth Kavuluru**

Division of Biomedical Informatics, Depts. of Biostatistics and Computer Science, University of Kentucky, Lexington, Kentucky

### Abstract

Building high accuracy text classifiers is an important task in biomedicine given the wealth of information hidden in unstructured narratives such as research articles and clinical documents. Due to large feature spaces, traditionally, discriminative approaches such as logistic regression and support vector machines with n-gram and semantic features (e.g., named entities) have been used for text classification where additional performance gains are typically made through feature selection and ensemble approaches. In this paper, we demonstrate that a more direct approach using convolutional neural networks (CNNs) outperforms several traditional approaches in biomedical text classification with the specific use-case of assigning medical subject headings (or MeSH terms) to biomedical articles. Trained annotators at the national library of medicine (NLM) assign on an average 13 codes to each biomedical article, thus semantically indexing scientific literature to support NLM's PubMed search system. Recent evidence suggests that effective automated efforts for MeSH term assignment start with binary classifiers for each term. In this paper, we use CNNs to build binary text classifiers and achieve an absolute improvement of over 3% in macro F-score over a set of selected hard-to-classify MeSH terms when compared with the best prior results on a public dataset. Additional experiments on 50 high frequency terms in the dataset also show improvements with CNNs. Our results indicate the strong potential of CNNs in biomedical text classification tasks.

### Keywords

convolutional neural networks; text classification; medical subject headings

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [Permissions@acm.org](mailto:Permissions@acm.org).

Correspondence to: Ramakanth Kavuluru.

## 1. Introduction

Text classification is an important problem with many applications in biomedicine. Specific problems such as identifying reportable cases of cancer from pathology reports, recognizing particular phenotypes from clinical narratives, determining the correct sense for the usage of an ambiguous word given its context (word sense disambiguation), and assigning medical subject headings (MeSH terms) to biomedical articles, can all be modeled as instances of the general text classification problem. Traditional approaches to text classification involve applying conditional models (e.g., logistic regression (LR), support vector machines(SVMs)) trained on features derived from the text including n-grams and named entities. Further performance gains are usually achieved with feature selection, data selection, and ensemble approaches such as voting, bagging, stacking, and boosting [40]. Deriving more interesting and relevant features based on the particular domain of interest and the nature of text are also popular additional enhancements, typically termed as feature engineering. For example, emoticons and hashtags form important new features for sentiment analysis of Twitter data [20]. In biomedicine identifying negated mentions of named entities and automatically deriving regular expression features from a set of domain specific seed patterns have been shown to be effective [7]. Exploiting definitional knowledge (concept glosses) and inter-concept relations from an external domain specific knowledge base such as the unified medical language system (UMLS) can also provide performance gains [33]. Overall, applying linear classifiers with domain specific feature engineering and ensemble modeling still produces state-of-the-art results for text classification.

The resurgence of deep neural networks (or deep nets) has paved ways to more general alternatives to supervised learning, especially in object classification. Deep nets obviate the laborious process of feature engineering and take upon the burden of automatically learning high level representations of input instances that are better suitable for the classification problem at hand. Deep nets have been initially applied to problems in computer vision but have been recently adapted to natural language processing (NLP) tasks [3,9,22] especially through learning distributed representations of textual segments (words, sentences, documents) as vectors in  $\mathbb{R}^d$ . These vectors directly guide primitive tasks such as part-of-speech tagging and statistical parsing as well as high level tasks such as text classification and machine translation. Convolution neural networks (CNNs) take advantage of the so called “convolutional filters” to automatically learn features that are more suitable to the task at hand. They have been actively used in biomedicine in image classification even before deep learning became popular in the recent past. However, CNNs and deep nets in general have not been explored until recently for text classification, and are currently used for sentiment/opinion mining [17,19,28] for short texts with fairly balanced class distributions. Our motivation in this paper is to evaluate CNNs on longer texts with highly skewed class distributions, a typical scenario encountered in biomedicine. We specifically demonstrate the potential of CNNs for learning better binary classifiers to assign medical subject headings (MeSH terms) to biomedical articles when compared with more conventional approaches.

Indexing biomedical articles with concepts from the controlled hierarchical MeSH vocabulary is an important task that has significant impact on how researchers search and retrieve relevant information. This is particularly essential given the exponential growth of

biomedical articles indexed by PubMed®, the main search system developed by the National Center for Biotechnology Information (NCBI). PubMed lets users search over 22 million biomedical citations available in the MEDLINE bibliographic database curated by the National Library of Medicine (NLM) from over 5000 leading biomedical journals in the world. To keep up with the explosion of information on various topics, users depend on queries involving MeSH terms that are assigned to each biomedical article. Once articles are indexed with MeSH terms, users can quickly search for articles that pertain to a specific subject of interest instead of relying solely on key word based searches. Besides this direct application, recent efforts also demonstrated that using the set of MeSH terms assigned to an article as its semantic proxy can be helpful in high level applications such as literature based knowledge discovery [8].

To mitigate indexing consistency issues and expedite the indexing process, there have been many recent efforts by researchers to develop automatic ways of assigning MeSH terms for indexing biomedical articles including efforts in the on-going BioASQ indexing challenge [32]. However, automated efforts (including ours) mostly focused on predicting MeSH terms for indexing based solely on the abstract and title text of the articles. This is because most full text articles are only available based on paid licenses not subscribed by many researchers. In this paper, we utilize recent advances in text classification using CNNs for assigning MeSH terms to biomedical articles based on the title and abstract text of the article. Jimeno-Yepes et al. [35] recently identified a set of nearly 29 hard-to-classify terms based on automatic indexing efforts at the NLM. They use a variety of classifiers in an ensemble setup and achieve better results than NLM's medical text indexer (MTI) program. We use their dataset and employ a CNN model proposed by Kim [19] to achieve an absolute improvement of over 3% in macro F-score over these selected terms. In a very recent attempt [15], Jimeno-Yepes et al. demonstrate the use of an extensive set of features to obtain improved results over 50 high frequency terms in the same dataset used in their 2013 paper. On these 50 terms, we achieve about 1% improvement in micro F-score and a comparable macro F-score. We believe that our results are the first to demonstrate the utility of CNNs for biomedical text classification especially for the scenarios with extreme class imbalance.

We discuss essential background on automated efforts to MeSH term assignment in Section 2 along with related efforts that use neural networks for natural language processing. In Section 3, we elaborate the specifics of the CNN model we use in our paper. We describe the dataset used for experiments, outline the methods compared, and specify the settings used for the CNNs in Section 4. We present and discuss our main results in Section 5 for three specific groups of MeSH terms. We also provide a brief analysis on what is actually being captured by the convolutions learned. In Section 6, we provide new results based on our recent experiments conducted for the 50 high frequency terms in the dataset. Section 7 summarizes our results and identifies some limitations of CNNs and directions for future work.

## 2. Background

Assigning MeSH terms to biomedical articles is an instance of the well known multi-label classification problem where multiple labels from the MeSH vocabulary need to be assigned to each input instance, which is a biomedical article. NLM initiated efforts in automatic MeSH term assignment with their MTI program that exploits terms from already tagged related citations in combination with named entity recognition (NER), unsupervised clustering, ad hoc indexing rules, and candidate term ranking heuristics in a pipeline [1]. MTI recommends MeSH terms for NLM indexers to assist in their efforts to expedite the indexing process. A few recent studies [13,18] apply the  $k$ -NN approach to obtain MeSH terms from a set of top  $k$ , already indexed, nearest neighbors and use the supervised learning-to-rank approach with novel features to rerank the MeSH terms from the neighbors. Other researchers considered different machine learning approaches with novel feature selection [36] and training data sampling [29] techniques.

A recent effort by Jimeno-Yepes et al. [14] uses a large dataset and meta-learning to train custom binary classifiers for each MeSH term and indexes the best performing model for each label to be applied on new abstracts. This method is typically known as the binary relevance approach where separate datasets of positive and negative examples are built for each label. Specifically, let  $T$  be the set of labels and let  $q = |T|$ . Binary relevance learns  $q$  binary classifiers, one for each label in  $T$ . It transforms the dataset into  $q$  separate datasets. For each label  $T_j$ , we obtain the dataset for the corresponding binary classifier by considering each document-label-set pair  $(D_i, \mathbf{Y}_i)$  and generating the document-label pair  $(D_i, T_j)$  when  $T_j \in \mathbf{Y}_i$  and generating the pair  $(D_i, \neg T_j)$  when  $T_j \notin \mathbf{Y}_i$ . The labels whose classifiers output a positive decision are finally considered the assigned labels. Recent state-of-the-art results [21] are obtained by combining the  $k$ -NN approach with the binary relevance approach, where the candidates from a few nearest neighbors are combined with top few predicted candidates from the binary relevance approach (based on classifier scores). This combined set of candidates is then ranked using learning-to-rank with a variety of features for each candidate that include neighborhood similarity scores from the  $k$ -NN approach and classifier scores from the binary relevance approach. Given this, we note that building high accuracy binary classifiers is crucial for MeSH term assignment and we use CNNs for this task in this paper.

Neural word representations have been shown to capture both semantic and syntactic information and a few recent approaches learn word vectors [3,9,22] (as elements of  $\mathbb{R}^d$ , where  $d$  is the dimension) in an unsupervised fashion from textual corpora. Deep nets have also been applied to sentence/document level classification problems especially sentiment analysis and opinion mining with relatively smaller and reasonably balanced datasets with few classes. Classification models were developed for sentiment analysis to take advantage of the structure of sentences. For instance in recursive neural networks [28], words of a sentence are recursively merged together using a nonlinear function until we have single vector. The elements of this vector are then used as features for classification. In our work, we build binary MeSH classifiers using the CNN model by Kim [19], which is closely related to time-delay neural networks and dynamic convolutional neural networks described by Kalchbrenner et. al [17]. Our main aims are to see if CNNs are helpful for longer

narratives (200 tokens per biomedical citation vs 20-50 tokens per opinion/sentiment containing narratives) and whether they can deal with large class imbalance for assigning certain MeSH terms. In the next section, we present the details of the CNN model.

### 3. CNN Model Details

The goal of the model described in this section is to build a binary classifier that outputs the probability estimates of unseen documents belonging to the positive class. This model is recently proposed by Kim [19]; we add additional inputs (see Section 6) to the softmax layer of this model besides traditional CNN features. In this section we furnish additional details of the model starting with a more intuitive explanation of CNNs and subsequently elaborating with a more detailed specification of the model and the training process.

The central concept in CNNs is the notion of *convolution filters* (CFs) that are traditionally used in signal processing. The general principle is to learn several CFs which are able to extract useful features from a document for the specific classification task based on the training dataset. This has proven to be very useful in computer vision [38] where convolutions learn high level features of an image (e.g., curves and faces). Before we go into the specifics, we provide a high-level overview of convolutions for text classification.

#### 3.1 Convolutions and CFs

A convolution is a binary operation involving the operands: a segment of text and a specific CF, both of which are represented as real matrices for our purposes. The output is a single real number. The matrix representation of the text segment, which is typically a contiguous sequence of words in the document, is composed of the word vectors of tokens that constitute it. A CF is also a matrix of the same dimensions as the text segment matrix. A specific CF operates on all contiguous segments of a document using a sliding window producing as many real number outputs as there are contiguous segments of a certain length in the document. This sequence of real numbers is called the feature map associated with the particular CF being used. Different CFs produce different feature maps, which can then be used as features in text classification. The overview of the convolution operation and CFs explained here forms the convolutional layer of the CNN. There is a conventional softmax layer that takes as input the feature maps and outputs class probability estimates. The main idea is to learn CFs (that is, the elements in the corresponding matrices) that give better feature maps that optimize our objective function. The learning process is based on predicting classes for training instances and making adjustments to the CF elements through back propagation of the gradients of the objective function (conditional log-likelihood of training data) being optimized. In this intuitive explanation, we left out many details including the mathematical definition of the convolution operation, the objective function, regularization (to deal with overfitting), and the stopping protocol for the learning process. The following section discusses these in detail.

#### 3.2 Model Specification

The architecture of the full CNN model used is shown in Figure 1. It has two layers including a single convolution layer and a fully connected output (softmax) layer.

The base component in the model is a word vector  $\mathbf{x} \in \mathbb{R}^d$ , where  $d$  is the dimension of the word vectors. A document is represented as a matrix  $D \in \mathbb{R}^{n \times d}$ , where  $n$  is the number of words in it and each row represents the word vector for the corresponding token. To simplify the equations in this section we will assume the ground truth for the document  $Y \in \mathbb{R}^2$  such that  $Y_2 = 1$  and  $Y_1 = 0$  ( $Y_2 = 0$  and  $Y_1 = 1$ ) when we are training on a positive (negative) instance. This is more aligned with the two output nodes of the final layer for the two corresponding classes (positive/negative) for each binary classifier. Although a single node would have been sufficient for binary classification, we chose to build our model with multiple nodes to simply have the code setup for multi-class classification for other text classification problems.

We define a CF  $\mathbf{W} \in \mathbb{R}^{h \times d}$ , where  $h$  is the number of words we wish the convolution filter to span, that is, the length of the sliding window. Let the 2-D convolution operation  $*$  be defined as

$$\mathbf{W} * D_{j:j+h-1} = \sum_{i=j}^{j+h-1} \sum_{k=0}^{d-1} \mathbf{W}_{i,k} D_{i,k}.$$

We next map a length  $h$  word window,  $D_{j:j+h-1}$ , of the document to a real number  $c_j \in \mathbb{R}$  using a non-linear function  $f$  as

$$c_j = f(\mathbf{W} * D_{j:j+h-1} + b), \quad (1)$$

where  $b \in \mathbb{R}$  represents the bias term. In this work  $f$  is a rectified linear function [12,23]. After convolving over the entire document using  $\mathbf{W}$ , we have the corresponding convolved feature map

$$\mathbf{c}(\mathbf{W}) = [c_1, c_2, \dots, c_{n-h+1}].$$

To overcome the issue of varying document lengths we perform max-pooling [10] operation

$$\hat{c}_{\mathbf{W}} = \max_i \mathbf{c}(\mathbf{W})_i,$$

which gives a single feature  $\hat{c}_{\mathbf{W}}$  corresponding to the feature map generated by  $\mathbf{W}$ . However, we will learn several CFs, say  $k$  of them,  $\mathbf{W}^1, \dots, \mathbf{W}^k$ , to create multiple feature maps leading to the corresponding single max-pooled features  $\hat{\mathbf{c}}_{\mathbf{W}^t}$ ,  $t = 1, \dots, k$ . These form a final max-pooled feature vector

$$\hat{\mathbf{c}}_{\mathcal{W}} = [\hat{c}_{\mathbf{W}^1}, \dots, \hat{c}_{\mathbf{W}^k}]^T,$$

where  $\mathcal{W} = \{\mathbf{W}^1, \dots, \mathbf{W}^k\}$ .

After obtaining  $\hat{\mathbf{c}}_{\mathcal{W}}$ , we add a final softmax layer. Let  $\mathbf{U} \in \mathbb{R}^{2 \times k}$  and  $b^U \in \mathbb{R}^2$  be the parameters of the softmax layer with weighted inputs

$$y_j = \mathbf{U}_j \hat{\mathbf{c}}_{\mathcal{W}} + b_j^U \quad (2)$$

and output label probability estimates

$$P(Y_j=1|D, \mathcal{W}, b, \mathbf{U}, b^U) = \frac{e^{y_j}}{\sum_i e^{y_i}}, \quad (3)$$

where  $\mathbf{U}_j$  is the  $j$ -th row of  $\mathbf{U}$ ,  $b_j^U$  is the  $j$ -th element of  $b^U$ , and  $Y_j$  is the  $j$ -th label for the document corresponding to matrix  $D$ .

If  $\mathcal{D}$  is the set of training document matrices, to learn each binary classifier we minimize

$$- \sum_{D \in \mathcal{D}} \log(P(Y_{\text{pos}}^D=1|D, \mathcal{W}, b, \mathbf{U}, b^U)), \quad (4)$$

where  $pos = 1$  ( $pos = 2$ ) if the corresponding document is a negative (positive) instance and  $Y^D$  is the ground truth for document represented by  $D$ . The parameters of the CNN ( $\mathcal{W}$ ,  $b$ ,  $\mathbf{U}$ ,  $b^U$  from equation 4 and the word vectors) that minimize this are obtained by calculating the gradient and using back propagation with the stochastic gradient descent approach. A subtle but crucial aspect that makes CNNs for NLP tasks different from those used in computer vision is that the base input components to the CNN, that is, the word vectors, are also modified using back propagation in addition to the traditional network weights. This is done by treating the word vector elements as network weights of the first and so called projection layer [27, Chapter 2] and modifying them just like any other network weight. However, the vector elements for a given word only change when the current instance contains that word, which happens often if it is a common word or if the dataset is large. We used the popular mini-batches [25] approach instead of updating parameters for each example. CNNs also warrant multiple epochs where the learning process goes through the entire training dataset multiple times in optimizing equation 4.

Instead of using the well known l1 or l2 regularization we use the dropout [30] option to prevent over-fitting during training. Specifically, instead of passing  $y_j$  (from equation 2) to the softmax function in equation 3 during training, we actually pass

$$\hat{y}_j = \mathbf{U}_j (\hat{\mathbf{c}}_{\mathcal{W}} \circ \mathbf{r}) + b_j^U,$$

where  $\circ$  refers to element-wise multiplication and  $\mathbf{r} \in \{0, 1\}^k$  is constructed with each  $r_j$  drawn from the Bernoulli distribution with parameter  $p$  (typically set to 0.5). Intuitively, this means that gradients are backpropagated only through unmasked elements where  $r_j = 1$ . During test time we scale the weights  $U$  such that

$$y_j = p \mathbf{U}_j \hat{\mathbf{c}}_{\mathcal{W}} + b_j^U.$$

This down weighting is essential since at training, on average, only half of the  $U$  edges were active, which is not true at test time. We also used early-stopping in order to help combat over-fitting. Typically early-stopping is done by simply terminating the training of the model when the desired score on a held out validation dataset does not increase in performance. However, we found this caused us to stop too early. To combat this we stopped training if there were 5 consecutive epochs in the training procedure that did not increase the validation score. We only saved the model on epochs that had an increase in F-score on the validation dataset. For example, if there was an increase in F-score in the second epoch on the validation dataset, and then training continued for five more epochs without any further increase, we kept the model from the 2nd epoch. Additional specifics on initialization of word vectors and CFs and other parameter choices are discussed in the next section.

## 4. Experimental Setup

In this section we discuss the dataset used in our paper along with the MeSH terms for which we built CNN binary classifiers. We also discuss CNN model parameter settings and other approaches compared with them. For all our experiments we used the Python based Theano [4] platform.

### 4.1 Dataset and MeSH Terms Used

We use a publicly available<sup>1</sup> dataset consisting of MED-LINE citations from November 2012 to February 2013. It contains 89,942 biomedical citations for training, 5000 for validation, and 48,911 for testing. The dataset, the MeSH terms for which the classifiers are built, and the methods compared with CNNs are those used in a prior effort by Jimeno-Yepes et. al [35] to facilitate direct comparison. We compare our model with 29 MeSH terms used in [35] split into three groups. The three groups were created based on how NLM's MTI program performs on each term: check tags, low recall terms, and low precision terms. Check-tags are a special set of popular MeSH terms (e.g, *humans*, *female*, and *adult*), which are typically checked for each article to be indexed. We built CNN models for the 12 popular check tags, the top 7 low recall terms, and the top 10 low precision terms based on MTI's judgments; so a total of 29 terms were considered based on the three groups in [35]. The goal was to see if the binary relevance approach based on supervised learning approaches can provide a better alternative to the  $k$ -NN based approach employed in MTI. In [35], the authors also conduct experiments on a fourth group of terms for which MTI had zero recall. These terms are very rare (about 1% of 48,911 citations in the test set) and given

<sup>1</sup>[http://ii.nlm.nih.gov/MTI\\_ML/index.shtml](http://ii.nlm.nih.gov/MTI_ML/index.shtml)



we use a validation dataset of only 5000 citations, we have excluded this group from our evaluation after preliminary results showed that there were not enough positive examples in the validation dataset for early stopping.

#### 4.2 CNN Parameters of Section 3.2

For all experiments, we used a word vector of size  $d = 300$ . Each word vector is initialized with values drawn uniformly from  $[-0.25, 0.25]$ . The CF  $\mathbf{W}$  values and softmax layer weights  $\mathbf{U}$  are drawn uniformly from  $[-0.1, 0.1]$ . We used three different CF sizes with window lengths  $h = 3, 4, \text{ and } 5$ . For each of these filter sizes we used 100 feature maps creating a total of 300 feature maps per classifier. The models were trained using AdaDelta [37], an adaptive learning rate method for stochastic gradient descent with a maximum of 15 epochs per classifier. We also used mini-batches of size 50 and we zero-padded the document at the beginning as needed. The dropout parameter  $p$  was set to 0.5. Some of these choices were based on settings used in the paper where this CNN model was first used [19] and others were based on our experimentation.

#### 4.3 Methods Compared

The following list has a total of 13 types of models we compare in our effort. The last four are based on CNN models and the the rest of them are from Jimeno-Yepes et. al [35].

- Naive Bayes (NB)
- Logistic Regression (LR)
- Support Vector Machines (SVM)
- Support Vector Machines with Huber Loss (SVM HL)
- AdaBoostM1 (Ada)
- AdaBoostM1 with Oversampling (Ada Over)
- Medical Text Indexer (MTI)
- Vote 2 – This predicts a positive label if any two of the preceding models make a positive prediction for any given example.
- Vote 3 – This is the same as Vote 2 except it requires that at least 3 of the base algorithms NB, LR, SVM, SVM HL, Ada, Ada Over, or MTI predict any given label.
- CNN-rand – This is our first CNN model which uses randomized initial word vectors and is trained as described in Section 3.2.
- CNN-pre – This model initializes word vectors with those obtaining from running Word2Vec [22] on all biomedical citations in PubMed in the last decade except for instances in the test set.
- CNN-Vote 2 – This model ensembles five CNN-rand models (with different word vector initializations) and one CNN-pre model for each term. Predictions are made just as in Vote 2, when at least two models make a positive prediction.

## 5. Main Results and Discussion

In this section we present and discuss results for each group of terms using models in Section 4.3. The F-scores of the models for the three MeSH term groups (from Section 4.1) are shown in Tables 1–3. In these tables the ‘positive’ column refers to the number of positive examples in the test dataset. The ‘prior-best’ column refers to the maximum F-score taken over all non-CNN models (from [35]) enumerated in Section 4.3 except MTI (we list MTI separately since it is our baseline). However, the non-CNN models that involve voting (Vote 2 and Vote 3) include MTI as a component of the ensemble.

### 5.1 CNN Model Performance

For check-tags results (Table 1), MTI is not included because for check-tags MTI has been modified to use AdaBoostM1 with oversampling instead of the  $k$ -NN method used for other terms. The check tags group has an average of 8701 positive examples per label in the test set and is a popular set of terms frequently used; most citations have at least one check tag. Our results for this group in Table 1 show that CNN models improve the F-score for all but one term in the group. Interestingly, a single CNN with pre-trained word vectors seems to perform better than using an ensemble of CNNs on very frequent terms such as Humans, Male, and Female. This is not unexpected since our voting scheme requires at least two models predicting a term. However, for less frequent check tags, ensemble models prove to be very useful. Finally, we also note that the our check tag scores in Table 1 also show considerable improvement over another prior effort that uses deep belief networks [34, Table 2].

The low precision group has an average of 1214 examples per term. The low precision/recall situations are common in unbalanced datasets due to the inherent complexity in the particular class being predicted and also owing to underfitting/overfitting issues in certain algorithms. However, in Table 2 voting with CNNs consistently outperforms other methods. Next, we look at the performance of the low recall group in Table 3. This group has an average of 2002 examples per label. Again we see that CNNs models perform better than other models except for one term for which MTI performs better. Finally, we consider the macro average scores for all models in each term group in Table 4 to assess their consistency. We can see that CNN ensembles provide the best results. Overall, CNNs seem to perform better than any other individual method. Ensembled CNNs perform better than any individual method and better than any ensemble method without CNNs.

The main difference between the CNN models we used and other methods we compared with, such as the Linear SVM and Logistic Regression, is that our method creates high level abstractions using CFs. However, even in cases when there is a small number of positive examples, CNN based models perform better than all other methods. In Table 4, CNN-Vote 2 is an ensemble containing only CNN models. However, even in the low precision group which has, on average, the least number of positive examples per term, CNN-Vote 2 still achieves 2% improvement over the best model that does not use CNNs. In all our voting models in this section, we purposefully avoided including MTI as a component model given the potential bias as the test set articles are already indexed in PubMed when we conducted our experiments. However, we believe using MTI as part of the voting ensembles for new

citations will yield better performance since MTI also incorporates nearest neighbor based predictions and as such offers complementary predictive power.

## 5.2 Visualizing Convolutions

Given the results are promising, we wanted to see what the CFs are actually capturing at a high level. For this we applied each of the CFs obtained after training the Humans model to all trigrams in the test dataset. We then identified the top five trigrams when ranked based on the first layer output score in equation 1. Table 5 shows four manually selected CFs of the 100 we used. CF-2 seems to have learned to capture information about small things, such as nanoparticles and cells. Raman images are actually an imaging/fingerprinting method for molecules. CF-4 seems to capture information from citations that are performing different types of analysis, which may have to do with numerical properties. These results are interesting in that instead of just learning to capture very specific important n-grams (like in logistic regression), CFs seem to capture different topics or aspects of the input citations to be classified.

## 6. Additional experiments with Frequent Terms

In a recent paper [15], Jimeno-Yepes et al. used the same dataset used in their prior effort and in our current effort to conduct a comprehensive feature engineering study over a set of 50 randomly selected high frequency terms from 63 terms that have a minimum frequency of 1500 citations in the dataset. They select a large set of feature types including named entities (and their hypernymic variants), journal, author, and author affiliations of the citation, noun phrases, and argumentative structure enhanced n-grams (based on methods, results, and discussion components of structured abstracts). They also treat title and abstract as separate fields and consider the corresponding features as having different types. They also include as features term sets from NLM's MTI indexing program and its components including the PubMed related citations (PRC) component that in a sense fetches nearest neighbors of the input citation. Besides the conventional SVMs with linear kernels, they also used  $SVM^{perf}$  [16] to directly optimize F-score (actually, a function that lower bounds F-score), which is a non-linear measure that cannot be directly optimized with traditional classifiers.

We conducted experiments with our CNN models over the same set of 50 high frequency terms (those used in [15]) whose results are shown in Table 6. In addition to the straightforward macro average discussed earlier, we also compute micro F-score over the 50 terms (where the contingency table counts of all classifier outputs are pooled to compute a single F-measure). As we can see, the single model with pre-trained word vectors without any of the feature engineering described earlier has a better micro F-score compared with the prior best score. However, the macro F-score is six points lower with the single pre-trained CNN. We note that the outer softmax layer in Figure 1 can take other features in addition to the CFs. Out of more than ten different types of features used in [15], we use two of them: PRC component based terms and named entities by running the named entity recognition tool MetaMap [2]. Both these are components of the MTI program. We did not use MTI predictions given the articles in the test set are already indexed in PubMed. Adding these

two features to the softmax layer improves the macro F-score over the pre-trained model by two percentage points with negligible loss in micro F-score.

From the vote-2 model 5 in Table 6, our macro F-score is lower by less than 1% compared with the best score produced by extensive feature engineering with a 2.5% increase in micro F-score over the corresponding feature engineered model that produced the best macro F-score. Model 6 is same as model 5 except voting is not used for the top four frequent terms, which is a natural choice given voting might decrease the performance of very frequent terms. This improves both micro and macro F-scores by 0.1% compared with model 5. Thus, our analysis shows that a single CNN based voting model (model 6) with two straightforward features (besides the CFs) almost achieves (macro score) or improves (micro score) over two separate feature engineered models that do not achieve comparable macro-micro score combinations.

## 7. Concluding Remarks

In this effort, we demonstrated that CNNs with CFs over word sequences are effective for biomedical text classification in achieving new state-of-the-art scores over more traditional linear classifiers especially when there is significant label imbalance. Specifically, using a well known public gold standard dataset, we achieved a macro F-score improvement of 3% over previous best scores on three specific groups of hard to classify MeSH terms (Table 4). Additional experiments over 50 high frequency terms also reveal that elaborate feature engineering can also be minimized with CNNs (Table 6) and that new features can be simply passed to the outer softmax layer of the CNN to achieve better results compared with using those features in traditional classifiers. We believe our results will further improve if we incorporate predictions from the MTI program given it offers a complementary approach to our methods. We, however, chose to exclude MTI terms from our experiments given potential bias that could occur with our particular test dataset. Overall, our results demonstrate the strength of CNNs in capturing high level features that lead to superior classification performance over other approaches that involve several well known linear classifiers with elaborate feature engineering and ensembling.

Next, we outline a few research directions we plan to explore in the immediate future. First, we will study deeper CNNs utilizing more sophisticated max-pooling procedures. In this effort we only have one convolutional layer followed by a softmax layer. It would be interesting to look at the types of higher order features CNNs generate at deeper layers for text classification. CNNs for image classification have been widely studied and have been shown to produce very interesting higher order features [38]. For text, topic models [5, 6, 31] have been extensively studied based on their ability to represent documents as a distributions of ‘topics’. Can deep models learn ‘topics’ or some other higher order representation to help with different tasks? In this paper, we only perform max-over-time pooling procedure after a CF operates on sliding windows of a chosen size. However, dynamic max-pooling methods have been proposed for text classification [17] with CNNs. Using a max-over-time method loses a significant amount of information, especially for long documents and we would like to explore these other alternatives.

A caveat of CNNs is that even with early stopping they typically take at least an order of magnitude more time than conventional classifiers. This is typically not an issue for building high quality classifiers for few labels. However, in multi-label classification scenarios over large label spaces, using CNNs as part of the binary relevance approach, which requires learning a different classifier per label, can be prohibitive. So we will look at methods for expanding CNNs to simultaneously predict over the entire label space of MeSH avoiding the binary relevance approach. We will explore different loss functions that are already proposed for general neural networks for multi-label classification [24, 39]. We would like to see if a deep CNN with a multi-label output can perform well on MeSH. Finally, we will explore methods that take advantage of label correlations and prior structure we know about the MeSH label space. While there have been several methods that take advantage of label correlations [26], few look at very large label spaces. There has been work in computer vision that exploits prior information about labels to improve prediction on classes that have very few examples [11]. We would like to explore whether they can be applied to a multi-label CNN for MeSH term assignment.

## Acknowledgments

We are grateful to an anonymous reviewer who correctly pointed out the potential for recall oriented bias in our results when we include NLM's MTI program in the voting ensemble. Our experiments show that the Web based MTI program resulted in high recall as the test instances are already indexed and are potentially being incorporated in MTI results. We have removed MTI from the voting ensembles and also from inputs to the softmax layer in Section 6. We updated the scores and performance improvement estimates accordingly, given this valid criticism. We also thank three other reviewers for their thorough assessment of our contribution and specific comments that helped improve the quality of this paper. This publication was supported by the National Center for Research Resources and the National Center for Advancing Translational Sciences, US National Institutes of Health (NIH), through Grant UL1TR000117. The content is solely the responsibility of the authors and does not necessarily represent the official views of the NIH.

## References

1. Aronson A, Mork J, Gay C, Humphrey S, Rogers W. The NLM indexing initiative's medical text indexer. Proceedings of MEDINFO. 2004
2. Aronson AR, Lang FM. An overview of metapmap: historical perspective and recent advances. Journal of the American Medical Informatics Association. 2010; 17(3):229–236. [PubMed: 20442139]
3. Bengio Y, Ducharme R, Vincent P, Janvin C. A neural probabilistic language model. The Journal of Machine Learning Research. 2003; 3:1137–1155.
4. Bergstra J, Breuleux O, Bastien F, Lamblin P, Pascanu R, Desjardins G, Turian J, Warde-Farley D, Bengio Y. Theano: a CPU and GPU math expression compiler. Proceedings of the Python for Scientific Computing Conference (SciPy). Jun.2010 Oral Presentation.
5. Blei D, Lafferty J. Correlated topic models. Advances in neural information processing systems. 2006; 18:147.
6. Blei DM, Ng AY, Jordan MI. Latent dirichlet allocation. the Journal of machine Learning research. 2003; 3:993–1022.
7. Bui DDA, Zeng-Treitler Q. Learning regular expressions for clinical text classification. Journal of the American Medical Informatics Association. 2014 amiajnl–2013.
8. Cameron D, Kavuluru R, Rindflesch TC, Sheth AP, Thirunarayan K, Bodenreider O. Context-driven automatic subgraph creation for literature-based discovery. Journal of biomedical informatics. 2015; 54:141–157. [PubMed: 25661592]

9. Collobert, R., Weston, J. Proceedings of the 25th international conference on Machine learning. ACM; 2008. A unified architecture for natural language processing: Deep neural networks with multitask learning; p. 160-167.
10. Collobert R, Weston J, Bottou L, Karlen M, Kavukcuoglu K, Kuksa P. Natural language processing (almost) from scratch. *Journal of Machine Learning Research*. 2011; 12:2493–2537.
11. Deng, J., Ding, N., Jia, Y., Frome, A., Murphy, K., Bengio, S., Li, Y., Neven, H., Adam, H. *Computer Vision–ECCV 2014*. Springer; 2014. Large-scale object classification using label relation graphs; p. 48-64.
12. Glorot X, Bordes A, Bengio Y. Deep sparse rectifier networks. *Proceedings of the 14th International Conference on Artificial Intelligence and Statistics JMLR W&CP Volume*. 2011; 15:315–323.
13. Huang M, Névéol A, Lu Z. Recommending mesh terms for annotating biomedical articles. *Journal of the American Medical Informatics Association*. 2011; 18(5):660–667. [PubMed: 21613640]
14. Jimeno-Yepes A, Mork JG, Demner-Fushman D, Aronson AR. A one-size-fits-all indexing method does not exist: Automatic selection based on meta-learning. *Journal of Computing Science and Engineering*. 2012; 6(2):151–160.
15. Jimeno Yepes A, Plaza L, Carrillo-de Albornoz J, Mork JG, Aronson AR. Feature engineering for medline citation categorization with mesh. *BMC Bioinformatics*. 2015; 16(1):113. [PubMed: 25887792]
16. Joachims, T. Proceedings of the 22nd international conference on Machine learning. ACM; 2005. A support vector method for multivariate performance measures; p. 377-384.
17. Kalchbrenner, N., Grefenstette, E., Blunsom, P. Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). Baltimore, Maryland: Association for Computational Linguistics; Jun. 2014 A convolutional neural network for modeling sentences; p. 655-665.
18. Kavuluru R, Lu Y. Leveraging output term co-occurrence frequencies and latent associations in predicting medical subject headings. *Data & Knowledge Engineering*. 2014; 94(Part B):189–201.
19. Kim, Y. Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP). Doha, Qatar: Association for Computational Linguistics; Oct. 2014 Convolutional neural networks for sentence classification; p. 1746-1751.
20. Kiritchenko S, Zhu X, Mohammad SM. Sentiment analysis of short informal texts. *Journal of Artificial Intelligence Research*. 2014:723–762.
21. Liu K, Wu J, Peng S, Zhai C, Zhu S. The fudan-uiuc participation in the BioASQ challenge task 2a: The antinomyra system. *Proceedings of Question Answering Lab at the Conference and Labs of the Evaluation Forum (CLEF)*. 2014
22. Mikolov T, Sutskever I, Chen K, Corrado GS, Dean J. Distributed representations of words and phrases and their compositionality. *Advances in Neural Information Processing Systems*. 2013:3111–3119.
23. Nair V, Hinton GE. Rectified linear units improve restricted Boltzmann machines. *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*. 2010:807–814.
24. Nam, J., Kim, J., Mencía, EL., Gurevych, I., Fürnkranz, J. *Machine Learning and Knowledge Discovery in Databases*. Springer; 2014. Large-scale multi-label text classification revisiting neural networks; p. 437-452.
25. Ngiam J, Coates A, Lahiri A, Prochnow B, Le QV, Ng AY. On optimization methods for deep learning. *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*. 2011:265–272.
26. Read J, Pfahringer B, Holmes G, Frank E. Classifier chains for multi-label classification. *Machine Learning*. 2011; 85(3):335–359.
27. Socher, R. PhD thesis. Department of Computer Science, Stanford University; 2014. Recursive Deep Learning for Natural Language Processing and Computer Vision.
28. Socher, R., Perelygin, A., Wu, JY., Chuang, J., Manning, CD., Ng, AY., Potts, C. Proceedings of the conference on empirical methods in natural language processing (EMNLP). Vol. 1631. Citeseer; 2013. Recursive deep models for semantic compositionality over a sentiment treebank; p. 1642

29. Sohn S, Kim W, Comeau DC, Wilbur WJ. Optimal training sets for bayesian prediction of MeSH assignment. *Journal of the American Medical Informatics Association*. 2008; 15(4):546–553. [PubMed: 18436913]
30. Srivastava N, Hinton G, Krizhevsky A, Sutskever I, Salakhutdinov R. Dropout: A simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*. 2014; 15(1): 1929–1958.
31. Steyvers, M., Smyth, P., Rosen-Zvi, M., Griffiths, T. Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining. ACM; 2004. Probabilistic author-topic models for information discovery; p. 306-315.
32. Tsatsaronis G, Balikas G, Malakasiotis P, Partalas I, Zschunke M, Alvers MR, Weissenborn D, Krithara A, Petridis S, Polychronopoulos D, et al. An overview of the bioasq large-scale biomedical semantic indexing and question answering competition. *BMC bioinformatics*. 2015; 16(1):138. [PubMed: 25925131]
33. Yepes AJ, Berlanga R. Knowledge based word-concept model estimation and refinement for biomedical text mining. *Journal of biomedical informatics*. 2015; 53:300–307. [PubMed: 25510606]
34. Yepes AJ, MacKinlay A, Bedo J, Garnavi R, Chen Q. Deep belief networks and biomedical text categorisation. *Proceedings of the Twelfth Annual Workshop of the Australasia Language Technology Association*. 2014:123.
35. Yepes, AJ., Mork, JG., Demner-Fushman, D., Aronson, AR. AMIA Annual Symposium Proceedings. Vol. 2013. American Medical Informatics Association; 2013. Comparison and combination of several MeSH indexing approaches; p. 709
36. Yetisgen-Yildiz, M., Pratt, W. Proceedings of AMIA Symposium. Vol. 2005. American Medical Informatics Association; 2005. The effect of feature representation on medline document classification; p. 849-853.
37. Zeiler MD. ADADELTA: an adaptive learning rate method. *arXiv preprint arXiv:1212.5701*. 2012
38. Zeiler, MD., Krishnan, D., Taylor, GW., Fergus, R. *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*. IEEE; 2010. Deconvolutional networks; p. 2528-2535.
39. Zhang ML, Zhou ZH. Multilabel neural networks with applications to functional genomics and text categorization. *Knowledge and Data Engineering, IEEE Transactions on*. 2006; 18(10):1338–1351.
40. Zhou, ZH. *Ensemble methods: foundations and algorithms*. CRC Press; 2012.

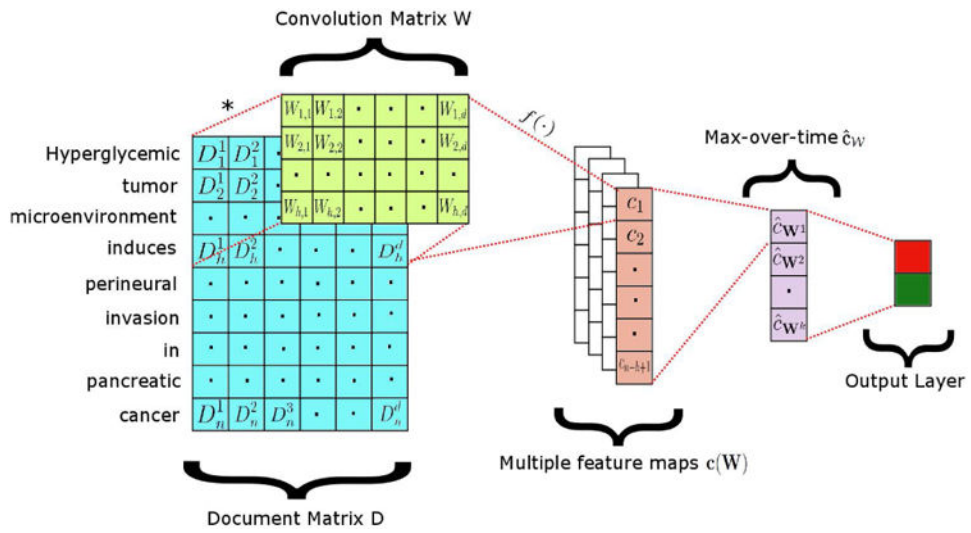


Figure 1. The CNN model layout

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript



Table 1

## The F-scores on test set for the check tags group

MeSH Term	Positive	Prior-best	CNN-rand	CNN-pre	CNN-Vote 2
Adolescent	3824	0.4708	0.5028	0.4951	<b>0.5335</b>
Adult	8792	0.6225	0.6593	0.6491	<b>0.6759</b>
Aged	6151	0.6005	0.6378	0.6427	<b>0.6516</b>
Aged, 80 and over	2328	0.3753	0.3909	0.3932	<b>0.4389</b>
Child, Preschool	1573	0.5129	0.4963	0.5296	0.5417
Female	16483	0.7764	0.7820	<b>0.7945</b>	0.7863
Humans	35967	0.9337	0.9322	<b>0.9407</b>	0.9339
Infant	1281	0.4796	0.4912	0.4926	<b>0.5378</b>
Male	15530	0.7582	0.7690	<b>0.7797</b>	0.7748
Middle Aged	8392	0.6731	0.7019	0.6911	<b>0.7058</b>
Swine	285	0.7323	0.7233	0.7220	0.7282
Young Adult	3807	0.3973	0.4007	0.4207	<b>0.4536</b>

Table 2

## The per-term F-scores for the low-precision MeSH term group

MeSH Term	Positive	MTI	Prior-best	CNN-rand	CNN-pre	CNN-Vote 2
Age Factors	889	0.0844	0.1748	0.2344	<b>0.2469</b>	0.2406
Brain	823	<b>0.5201</b>	0.4700	0.4640	0.4472	0.5014
Cell Line	781	0.2876	<b>0.3059</b>	0.2243	0.2847	0.2989
Cells, Cultured	1079	0.3046	0.3894	0.2893	0.3347	<b>0.3979</b>
Models, Molecular	851	0.4292	0.4763	0.3764	0.3740	<b>0.4914</b>
Molecular Sequence Data	1527	<b>0.5495</b>	0.5118	0.4047	0.4532	0.5178
RNA, Messenger	628	0.4477	0.4626	0.4675	0.3782	<b>0.4743</b>
Severity of Illness Index	751	0.1824	0.2415	0.2299	0.2152	<b>0.2489</b>
Time Factors	2153	0.0980	0.1513	0.1608	0.1574	<b>0.2041</b>
United States	2658	0.3585	0.4128	0.3956	0.4693	<b>0.4613</b>

**Table 3**  
**The per-term F-scores for the low-recall MeSH term group**

Mesh Term	Positive	MTI	prior-best	CNN-rand	CNN-pre	CNN-Vote 2
Child	2780	0.5836	0.5854	0.6202	0.6154	<b>0.6410</b>
Follow-Up Studies	1470	0.0407	0.2741	0.2989	0.3337	<b>0.3382</b>
Reproducibility of Results	1206	0.3191	0.3722	0.3316	0.3344	<b>0.3923</b>
Retrospective Studies	2183	0.6608	0.6592	0.6695	0.6681	<b>0.6756</b>
Risk Assessment	1014	0.2556	0.2189	0.2369	0.2252	<b>0.2661</b>
Risk Factors	2365	<b>0.4989</b>	0.4774	0.4748	0.4716	0.4878
Treatment Outcome	2999	0.4202	0.4421	0.3875	0.3985	<b>0.4553</b>

**Table 4**  
**The macro-averaged F-scores for each individual method for all three groups**

Model	Check Tags	Low P	Low R
MTI	–	0.3262	0.3970
NB	0.4191	0.1504	0.2588
LR	0.5441	0.2682	0.3653
SVM	0.5549	0.2606	0.3647
SVM HL	0.5437	0.2654	0.3587
Ada	0.5237	0.1720	0.3398
Ada Over	0.5646	0.2764	0.3819
Vote 2	0.6072	0.3596	0.4315
Vote 3	0.5885	0.2825	0.3933
CNN-rand	0.6240	0.3247	0.4313
CNN-pre	0.6293	0.3361	0.4353
CNN-Vote 2	<b>0.6469</b>	<b>0.3837</b>	<b>0.4652</b>

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

**Table 5**  
**Top five trigram activations from four different convolutions from the Humans check tag model**

Convolution Filter 1	Convolution Filter 2	Convolution Filter 3	Convolution Filter 4
generate strong evidence	antibody conjugated nanoparticles	variety of both	mathematical model
academic performance in	nanoparticles on intact	temperament and character	analysis thus we
of chemical reactions	to raman images	studies revealed survival	like climate ph
theory and experiment	of cells and	electrochemical biosensor was	of tumor uptake
techniques and vacuum	membranes where the	features of individuals	benthic cyanobacterial mats

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

**Table 6**  
**The micro and macro averaged F-scores for top 50 high frequency terms**

Model ID	Model	Micro-F	Macro-F
1	Jimeno-Yepes Best Micro (SVMLight +MTI)	0.7135	0.5128
2	Jimeno-Yepes Best Macro (SVM <sup>perf</sup> with all features)	0.6922	<b>0.5565</b>
3	CNN-pre	0.7175	0.4963
4	CNN-pre+MetaMap+PRC	0.7170	0.5185
5	CNN-Vote 2 with Model 4	0.7200	0.5473
6	Model 5 and best single models for top four terms	<b>0.7217</b>	0.5482

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript