

# Semi-supervised Clustering for Short Text via Deep Representation Learning

Zhiguo Wang and Haitao Mi and Abraham Ittycheriah

IBM T.J. Watson Research Center

Yorktown Heights, NY, USA

{zhigwang, hmi, abei}@us.ibm.com

## Abstract

In this work, we propose a semi-supervised method for short text clustering, where we represent texts as distributed vectors with neural networks, and use a small amount of labeled data to specify our intention for clustering. We design a novel objective to combine the representation learning process and the k-means clustering process together, and optimize the objective with both labeled data and unlabeled data iteratively until convergence through three steps: (1) assign each short text to its nearest centroid based on its representation from the current neural networks; (2) re-estimate the cluster centroids based on cluster assignments from step (1); (3) update neural networks according to the objective by keeping centroids and cluster assignments fixed. Experimental results on four datasets show that our method works significantly better than several other text clustering methods.

## 1 Introduction

Text clustering is a fundamental problem in text mining and information retrieval. Its task is to group similar texts together such that texts within a cluster are more similar to texts in other clusters. Usually, a text is represented as a bag-of-words or term frequency-inverse document frequency (TF-IDF) vector, and then the k-means algorithm (MacQueen, 1967) is performed to partition a set of texts into homogeneous groups.

However, when dealing with short texts, the characteristics of short text and clustering task raise several issues for the conventional unsupervised clustering algorithms. First, the number of unique words in each short text is small, as a re-

- |                                                                                                                                                                                                       |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| (a) What's the color of apples?<br>(b) When will this apple be ripe?<br>(c) Do you like apples?<br>(d) What's the color of oranges?<br>(e) When will this orange be ripe?<br>(f) Do you like oranges? |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

Table 1: Examples for short text clustering.

sult, the lexical sparsity issue usually leads to poor clustering quality (Dhillon and Guan, 2003). Second, for a specific short text clustering task, we have prior knowledge or particular intentions before clustering, while fully unsupervised approaches may learn some classes the other way around. Take the sentences in Table 1 for example, those sentences can be clustered into different partitions based on different intentions: *apple* {a, b, c} and *orange* {d, e, f} with a fruit type intention, or *what-question* {a, d}, *when-question* {b, e}, and *yes/no-question* cluster {c, f} with a question type intension.

To address the lexical sparsity issue, one direction is to enrich text representations by extracting features and relations from Wikipedia (Banerjee et al., 2007) or an ontology (Fodeh et al., 2011). But this approach requires the annotated knowledge, which is also language dependent. So the other direction, which directly encode texts into distributed vectors with neural networks (Hinton and Salakhutdinov, 2006; Xu et al., 2015), becomes more interesting. To tackle the second problem, semi-supervised approaches (e.g. (Bilenko et al., 2004; Davidson and Basu, 2007; Bair, 2013)) have gained significant popularity in the past decades. Our question is can we have a unified model to integrate neural networks into the semi-supervised framework?

In this paper, we propose a unified framework for the short text clustering task. We employ a

deep neural network model to represent short sentences, and integrate it into a semi-supervised algorithm. Concretely, we extend the objective in the classical unsupervised k-means algorithm by adding a penalty term from labeled data. Thus, the new objective covers three key groups of parameters: centroids of clusters, the cluster assignment for each text, and the parameters within deep neural networks. In the training procedure, we start from random initialization of centroids and neural networks, and then optimize the objective iteratively through three steps until converge:

- (1) assign each short text to its nearest centroid based on its representation from the current neural networks;
- (2) re-estimate cluster centroids based on cluster assignments from step (1);
- (3) update neural networks according to the objective by keeping centroids and cluster assignments fixed.

Experimental results on four different datasets show that our method achieves significant improvements over several other text clustering methods.

In following parts, we first describe our neural network models for text representation (Section 2). Then we introduce our semi-supervised clustering method and the learning algorithm (Section 3). Finally, we evaluate our method on four different datasets (Section 4).

## 2 Representation Learning for Short Texts

We represent each word with a dense vector  $w$ , so that a short text  $s$  is first represented as a matrix  $S = [w_1, \dots, w_{|s|}]$ , which is a concatenation of all vectors of  $w$  in  $s$ ,  $|s|$  is the length of  $s$ . Then we design two different types of neural networks to ingest the word vector sequence  $S$ : the convolutional neural networks (CNN) and the long short-term memory (LSTM). More formally, we define the presentation function as  $x = f(s)$ , where  $x$  represents the vector of the text  $s$ . We test two encoding functions (CNN and LSTM) in our experiments.

Inspired from Kim (2014), our CNN model views the sequence of word vectors as a matrix, and applies two sequential operations: *convolution* and *max-pooling*. Then, a fully connected layer is

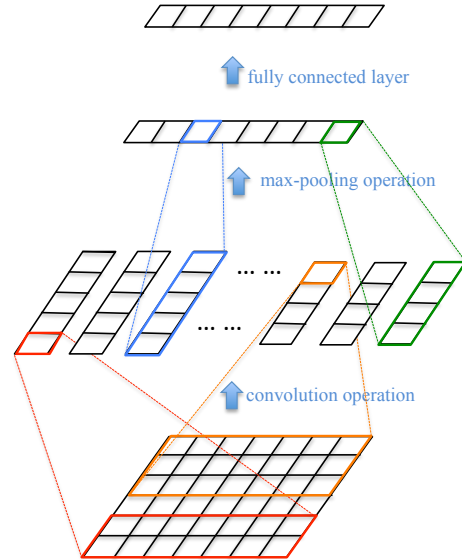


Figure 1: CNN for text representation learning.




Figure 2: LSTM for text representation learning.

employed to convert the final representation vector into a fixed size. Figure 1 gives the diagram of the CNN model. In the *convolution* operation, we define a list of filters  $\{w_o\}$ , where the shape of each filter is  $d \times h$ ,  $d$  is the dimension of word vectors and  $h$  is the window size. Each filter is applied to a patch (a window size  $h$  of vectors) of  $S$ , and generates a feature. We apply this filter to all possible patches in  $S$ , and produce a series of features. The number of features depends on the shape of the filter  $w_o$  and the length of the input short text. To deal with variable feature size, we perform a *max-pooling* operation over all the features to select the maximum value. Therefore, after the two operations, each filter generates only one feature. We define several filters by varying the window size and the initial values. Thus, a vector of features is captured after the *max-pooling* operation, and the feature dimension is equal to the number of filters.

Figure 2 gives the diagram of our LSTM model. We implement the standard LSTM block described in Graves (2012). Each word vector is fed into the LSTM model sequentially, and the

mean of the hidden states over the entire sentence is taken as the final representation vector.

### 3 Semi-supervised Clustering for Short Texts

#### 3.1 Revisiting K-means Clustering

Given a set of texts  $\{s_1, s_2, \dots, s_N\}$ , we represent them as a set of data points  $\{x_1, x_2, \dots, x_N\}$ , where  $x_i$  can be a bag-of-words or TF-IDF vector in traditional approaches, or a dense vector in Section 2. The task of text clustering is to partition the data set into some number  $K$  of clusters, such that the sum of the squared distance of each data point to its closest cluster centroid is minimized. For each data point  $x_n$ , we define a set of binary variables  $r_{nk} \in \{0, 1\}$ , where  $k \in \{1, \dots, K\}$  describing which of the  $K$  clusters  $x_n$  is assigned to. So that if  $x_n$  is assigned to cluster  $k$ , then  $r_{nk} = 1$ , and  $r_{nj} = 0$  for  $j \neq k$ . Let's define  $\mu_k$  as the centroid of the  $k$ -th cluster. We can then formulate the objective function as

$$J_{unsup} = \sum_{n=1}^N \sum_{k=1}^K r_{nk} \|x_n - \mu_k\|^2 \quad (1)$$

Our goal is to find the values of  $\{r_{nk}\}$  and  $\{\mu_k\}$  so as to minimize  $J_{unsup}$ .

The k-means algorithm optimizes  $J_{unsup}$  through the gradient descent approach, and results in an iterative procedure (Bishop, 2006). Each iteration involves two steps: *E-step* and *M-step*. In the *E-step*, the algorithm minimizes  $J_{unsup}$  with respect to  $\{r_{nk}\}$  by keeping  $\{\mu_k\}$  fixed.  $J_{unsup}$  is a linear function for  $\{r_{nk}\}$ , so we can optimize for each data point separately by simply assigning the  $n$ -th data point to the closest cluster centroid. In the *M-step*, the algorithm minimizes  $J_{unsup}$  with respect to  $\{\mu_k\}$  by keeping  $\{r_{nk}\}$  fixed.  $J_{unsup}$  is a quadratic function of  $\{\mu_k\}$ , and it can be minimized by setting its derivative with respect to  $\{\mu_k\}$  to zero.

$$\frac{\partial J_{unsup}}{\partial \mu_k} = 2 \sum_{n=1}^N r_{nk} (x_n - \mu_k) = 0 \quad (2)$$

Then, we can easily solve  $\{\mu_k\}$  as

$$\mu_k = \frac{\sum_{n=1}^N r_{nk} x_n}{\sum_{n=1}^N r_{nk}} \quad (3)$$

In other words,  $\mu_k$  is equal to the mean of all the data points assigned to cluster  $k$ .

#### 3.2 Semi-supervised K-means with Neural Networks

The classical k-means algorithm only uses unlabeled data, and solves the clustering problem under the unsupervised learning framework. As already mentioned, the clustering results may not be consistent to our intention. In order to acquire useful clustering results, some supervised information should be introduced into the learning procedure. To this end, we employ a small amount of labeled data to guide the clustering process.

Following Section 2, we represent each text  $s$  as a dense vector  $x$  via neural networks  $f(s)$ . Instead of training the text representation model separately, we integrate the training process into the k-means algorithm, so that both the labeled data and the unlabeled data can be used for representation learning and text clustering. Let us denote the labeled data set as  $\{(s_1, y_1), (s_2, y_2), \dots, (s_L, y_L)\}$ , and the unlabeled data set as  $\{s_{L+1}, s_{L+2}, \dots, s_N\}$ , where  $y_i$  is the given label for  $s_i$ . We then define the objective function as:

$$\begin{aligned} J_{semi} = & \alpha \sum_{n=1}^N \sum_{k=1}^K r_{nk} \|f(s_n) - \mu_k\|^2 \\ & + (1 - \alpha) \sum_{n=1}^L \{ \|f(s_n) - \mu_{g_n}\|^2 + \\ & \sum_{j \neq g_n} [l + \|f(s_n) - \mu_{g_n}\|^2 - \|f(s_n) - \mu_j\|^2]_+ \} \end{aligned} \quad (4)$$

The objective function contains two terms. The first term is adapted from the unsupervised k-means algorithm in Eq. (1), and the second term is defined to encourage labeled data being clustered in correlation with the given labels.  $\alpha \in [0, 1]$  is used to tune the importance of unlabeled data. The second term contains two parts. The first part penalizes large distance between each labeled instance and its correct cluster centroid, where  $g_n = G(y_n)$  is the cluster ID mapped from the given label  $y_n$ , and the mapping function  $G(\cdot)$  is implemented with the Hungarian algorithm (Munkres, 1957). The second part is denoted as a hinge loss with a margin  $l$ , where  $[x]_+ = \max(x, 0)$ . This part incurs some loss if the distance to the correct centroid is not shorter (by the margin  $l$ ) than distances to any of incorrect cluster centroids.

There are three groups of parameters in  $J_{semi}$ : the cluster assignment of each text  $\{r_{nk}\}$ , the clus-

- |                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <ol style="list-style-type: none"> <li>1. Initialize <math>\{\mu_k\}</math> and <math>f(\cdot)</math>.</li> <li>2. <i>assign_cluster</i>: Assign each text to its nearest cluster centroid.</li> <li>3. <i>estimate_centroid</i>: Estimate the cluster centroids based on the cluster assignments from step 2.</li> <li>4. <i>update_parameter</i>: Update parameters in neural networks.</li> <li>5. Repeat step 2 to 4 until convergence.</li> </ol> |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

Table 2: Pseudocode for semi-supervised clustering

ter centroids  $\{\mu_k\}$ , and the parameters within the neural network model  $f(\cdot)$ . Our goal is to find the values of  $\{r_{nk}\}$ ,  $\{\mu_k\}$  and parameters in  $f(\cdot)$ , so as to minimize  $J_{semi}$ . Inspired from the k-means algorithm, we design an algorithm to successively minimize  $J_{semi}$  with respect to  $\{r_{nk}\}$ ,  $\{\mu_k\}$ , and parameters in  $f(\cdot)$ . Table 2 gives the corresponding pseudocode. First, we initialize the cluster centroids  $\{\mu_k\}$  with the k-means++ strategy (Arthur and Vassilvitskii, 2007), and randomly initialize all the parameters in the neural network model. Then, the algorithm iteratively goes through three steps (*assign\_cluster*, *estimate\_centroid*, and *update\_parameter*) until  $J_{semi}$  converges.

The *assign\_cluster* step minimizes  $J_{semi}$  with respect to  $\{r_{nk}\}$  by keeping  $f(\cdot)$  and  $\{\mu_k\}$  fixed. Its goal is to assign a cluster ID for each data point. We can see that the second term in Eq. (4) has no relation with  $\{r_{nk}\}$ . Thus, we only need to minimize the first term by assigning each text to its nearest cluster centroid, which is identical to the *E-step* in the k-means algorithm. In this step, we also calculate the mappings between the given labels  $\{y_i\}$  and the cluster IDs (with the Hungarian algorithm) based on cluster assignments of all labeled data.

The *estimate\_centroid* step minimizes  $J_{semi}$  with respect to  $\{\mu_k\}$  by keeping  $\{r_{nk}\}$  and  $f(\cdot)$  fixed, which corresponds to the *M-step* in the k-means algorithm. It aims to estimate the cluster centroids  $\{\mu_k\}$  based on the cluster assignments  $\{r_{nk}\}$  from the *assign\_cluster* step. The second term in Eq. (4) makes each labeled instance involved in the estimating process of cluster centroids. By solving  $\partial J_{semi} / \partial \mu_k = 0$ , we get

$$\mu_k = \frac{\sum_{n=1}^N \alpha r_{nk} f(s_n) + \sum_{n=1}^L w_{nk} f(s_n)}{\sum_{n=1}^N \alpha r_{nk} + \sum_{n=1}^L w_{nk}} \quad (5)$$

$$w_{nk} = (1 - \alpha) (I'_{nk} + \sum_{j \neq g_n} I''_{nkj} - \sum_{j \neq g_n} I'''_{nkj})$$

$$I'_{nk} = \delta(k, g_n)$$

$$I''_{nkj} = \delta(k, j) \cdot \delta'_{nj}$$

$$I'''_{nkj} = (1 - \delta(k, j)) \cdot \delta'_{nj}$$

$$\delta'_{nj} = \delta(l + \|f(s_n) - \mu_{g_n}\|^2 - \|f(s_n) - \mu_j\|^2 > 0) \quad (6)$$

where  $\delta(x_1, x_2) = 1$  if  $x_1$  is equal to  $x_2$ , otherwise  $\delta(x_1, x_2) = 0$ , and  $\delta(x) = 1$  if  $x$  is true, otherwise  $\delta(x) = 0$ . The first term in the numerator of Eq. (5) is the contributions from all data points, and  $\alpha r_{nk}$  is the weight of  $s_n$  for  $\mu_k$ . The second term is acquired from labeled data, and  $w_{nk}$  is the weight of a labeled instance  $s_n$  for  $\mu_k$ .

The *update\_parameter* step minimizes  $J_{semi}$  with respect to  $f(\cdot)$  by keeping  $\{r_{nk}\}$  and  $\{\mu_k\}$  fixed, which has no counterpart in the k-means algorithm. The main goal is to update parameters for the text representation model. We take  $J_{semi}$  as the loss function, and train neural networks with the *Adam* algorithm (Kingma and Ba, 2014).

## 4 Experiment

### 4.1 Experimental Setting

We evaluate our method on four short text datasets. (1) *question\_type* is the TREC question dataset (Li and Roth, 2002), where all the questions are classified into 6 categories: abbreviation, description, entity, human, location and numeric. (2) *ag\_news* dataset contains short texts extracted from the AG’s news corpus, where all the texts are classified into 4 categories: World, Sports, Business, and Sci/Tech (Zhang and LeCun, 2015). (3) *dbpedia* is the DBpedia ontology dataset, which is constructed by picking 14 non-overlapping classes from DBpedia 2014 (Lehmann et al., 2014). (4) *yahoo\_answer* is the 10 topics classification dataset extracted from Yahoo! Answers Comprehensive Questions and Answers version 1.0 dataset by Zhang and LeCun

dataset	class#	total#	labeled#
question_type	6	5,953	595
ag_news	4	4,000	400
dbpedia	14	14,000	1,400
yahoo_answer	10	10,000	1,000

Table 3: Statistics for the short text datasets

(2015). We use all the 5,952 questions for the *question\_type* dataset. But the other three datasets contain too many instances (e.g. 1,400,000 instances in *yahoo\_answer*). Running clustering experiments on such a large dataset is quite inefficient. Following the same solution in (Xu et al., 2015), we randomly choose 1,000 samples for each classes individually for the other three datasets. Within each dataset, we randomly sample 10% of the instances as labeled data, and evaluate the performance on the remaining 90% instances<sup>1</sup>. Table 3 summarizes the statistics of these datasets.

In all experiments, we set the size of word vector dimension as  $d=300$ <sup>2</sup>, and pre-train the word vectors with the *word2vec* toolkit (Mikolov et al., 2013) on the English Gigaword (LDC2011T07). The number of clusters is set to be the same number of labels in the dataset. The clustering performance is evaluated with two metrics: Adjusted Mutual Information (AMI) (Vinh et al., 2009) and accuracy (ACC) (Amigó et al., 2009). In order to show the statistical significance, the performance of each experiment is the average of 10 trials.

## 4.2 Model Properties

There are several hyper-parameters in our model, e.g., the output dimension of the text representation models, and the  $\alpha$  in Eq. (4). The choice of these hyper-parameters may affect the final performance. In this subsection, we present some experiments to demonstrate the properties of our model, and find a good configuration that we use to evaluate our final model. All the experiments in this

<sup>1</sup>We didn't split dataset into train/dev/test portions which is commonly used for classification tasks, because it is not the convention for clustering task. First, the goal of clustering is to group given instances into clusters, instead of applying the trained model to new instances. Second, the evaluation process requires the clustering result of the whole set to map the clustering labels to the annotated labels.

<sup>2</sup>We tuned different dimensions for word vectors. When the size is small (50 or 100), performance drops significantly. When the size is larger (300, 500 or 1000), the curve flattens out. To make our model more efficient, we fixed it as 300.




Figure 3: Influence of the short text representation model, where the x-axis is the output dimension of the text representation models.




Figure 4: Influence of unlabeled data, where the x-axis is  $\alpha$  in Eq. (4).

subsection were performed on the *question\_type* dataset.

First, we evaluated the effectiveness of the output dimension in text representation models. We switched the dimension size among  $\{50, 100, 300, 500, 1000\}$ , and fixed the other options as:  $\alpha = 0.5$ , the filter types in the CNN model including  $\{\text{unigram, bigram, trigram}\}$  and 500 filters for each type. Figure 3 presents the AMIs from both CNN and LSTM models. We found that 100 is the best output dimension for both CNN and LSTM models. Therefore, we set the output dimension as 100 in the following experiments.

Second, we studied the effect of  $\alpha$  in Eq. (4), which tunes the importance of unlabeled data. We varied  $\alpha$  among  $\{0.00001, 0.0001, 0.001, 0.01, 0.1\}$ , and remain the other options as the last experiment. Figure 4 shows the AMIs from both CNN and LSTM models. We found that the clustering performance is not good when using a very small  $\alpha$ . By increasing the value of  $\alpha$ , we acquired progressive improvements, and reached to the peak point at  $\alpha=0.01$ . After that, the performance dropped. Therefore, we choose  $\alpha=0.01$  in




Figure 5: Influence of labeled data, where the x-axis is the ratio of data with given labels.

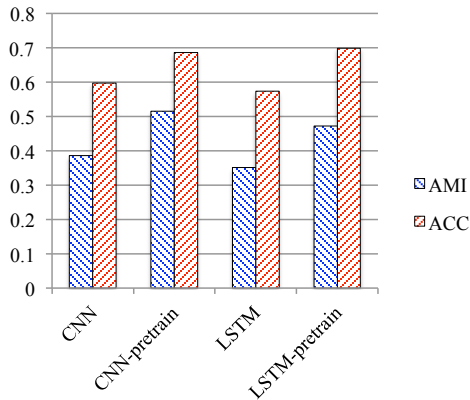


Figure 6: Influence of the pre-training strategy.

the following experiments. This results also indicate that the unlabeled data are useful for the text representation learning process.

Third, we tested the influence of the size of labeled data. We tuned the ratio of labeled instances from the whole dataset among [1%, 10%], and kept the other configurations as the previous experiment. The AMIs are shown in Figure 5. We can see that the more labeled data we use, the better performance we get. Therefore, the labeled data are quite useful for the clustering process.

Fourth, we checked the effect of the pre-training strategy for our models. We added a softmax layer on top of our CNN and LSTM models, where the size of the output layer is equal to the number of labels in the dataset. We then trained the model through the classification task using all labeled data. After this process, we removed the top layer, and used the remaining parameters to initialize our CNN and LSTM models. The performance for our models with and without pre-training strategy are given in Figure 6. We can see that the pre-training

strategy is quite effective for our models. Therefore, we use the pre-training strategy in the following experiments.

### 4.3 Comparing with other Models

In this subsection, we compared our method with some representative systems. We implemented a series of clustering systems. All of these systems are based on the k-means algorithm, but they represent short texts differently:

**bow** represents each text as a bag-of-words vector.

**tf-idf** represents each text as a TF-IDF vector.

**average-vec** represents each text with the average of all word vectors within the text.

**metric-learn-bow** employs the metric learning method proposed by Weinberger et al. (2005), and learns to project a bag-of-words vector into a 300-dimensional vector based on labeled data.

**metric-learn-idf** uses the same metric learning method, and learns to map a TF-IDF vector into a 300-dimensional vector based on labeled data.

**metric-learn-ave-vec** also uses the metric learning method, and learns to project an averaged word vector into a 100-dimensional vector based on labeled data.

We designed two classifiers (cnn-classifier and lstm-classifier) by adding a softmax layer on top of our CNN and LSTM models. We trained these two classifiers with labeled data, and utilized them to predict labels for unlabeled data. We also built two text representation models (“cnn-represent.” and “lstm-represent.”) by setting parameters of our CNN and LSTM models with the corresponding parameters in cnn-classifier and lstm-classifier. Then, we used them to represent short texts into vectors, and applied the k-means algorithm for clustering.

Table 4 summarizes the results of all systems on each dataset, where “semi-cnn” is our semi-supervised clustering algorithm with the CNN model, and “semi-lstm” is our semi-supervised clustering algorithm with the LSTM model. We



		question_type		ag_news		dbpedia		yahoo_answer	
		AMI	ACC	AMI	ACC	AMI	ACC	AMI	ACC
Unsup.	bow	0.028	0.257	0.029	0.311	0.578	0.546	0.019	0.140
	tf-idf	0.031	0.259	0.168	0.449	0.558	0.527	0.023	0.145
	average-vec	0.135	0.356	0.457	0.737	0.610	0.619	0.077	0.222
Sup.	metric-learn-bow	0.104	0.380	0.459	0.776	0.808	0.854	0.125	0.329
	metric-learn-idf	0.114	0.379	0.443	0.765	0.821	0.876	0.150	0.368
	metric-learn-ave-vec	0.304	0.553	0.606	0.851	0.829	0.879	0.221	0.400
	cnn-classifier	0.511	<b>0.771</b>	0.554	0.771	0.879	0.938	0.285	0.501
	cnn-represent.	0.442	0.618	0.604	0.833	0.864	0.899	0.210	0.334
	lstm-classifier	0.482	0.741	0.524	0.763	0.862	0.928	0.283	0.512
	lstm-represent.	0.421	0.618	0.535	0.771	0.667	0.706	0.152	0.272
Semisup.	semi-cnn	<b>0.529</b>	0.739	<b>0.662</b>	<b>0.876</b>	<b>0.894</b>	<b>0.945</b>	<b>0.338</b>	<b>0.554</b>
	semi-lstm	0.492	0.712	0.599	0.830	0.788	0.802	0.187	0.337

Table 4: Performance of all systems on each dataset.

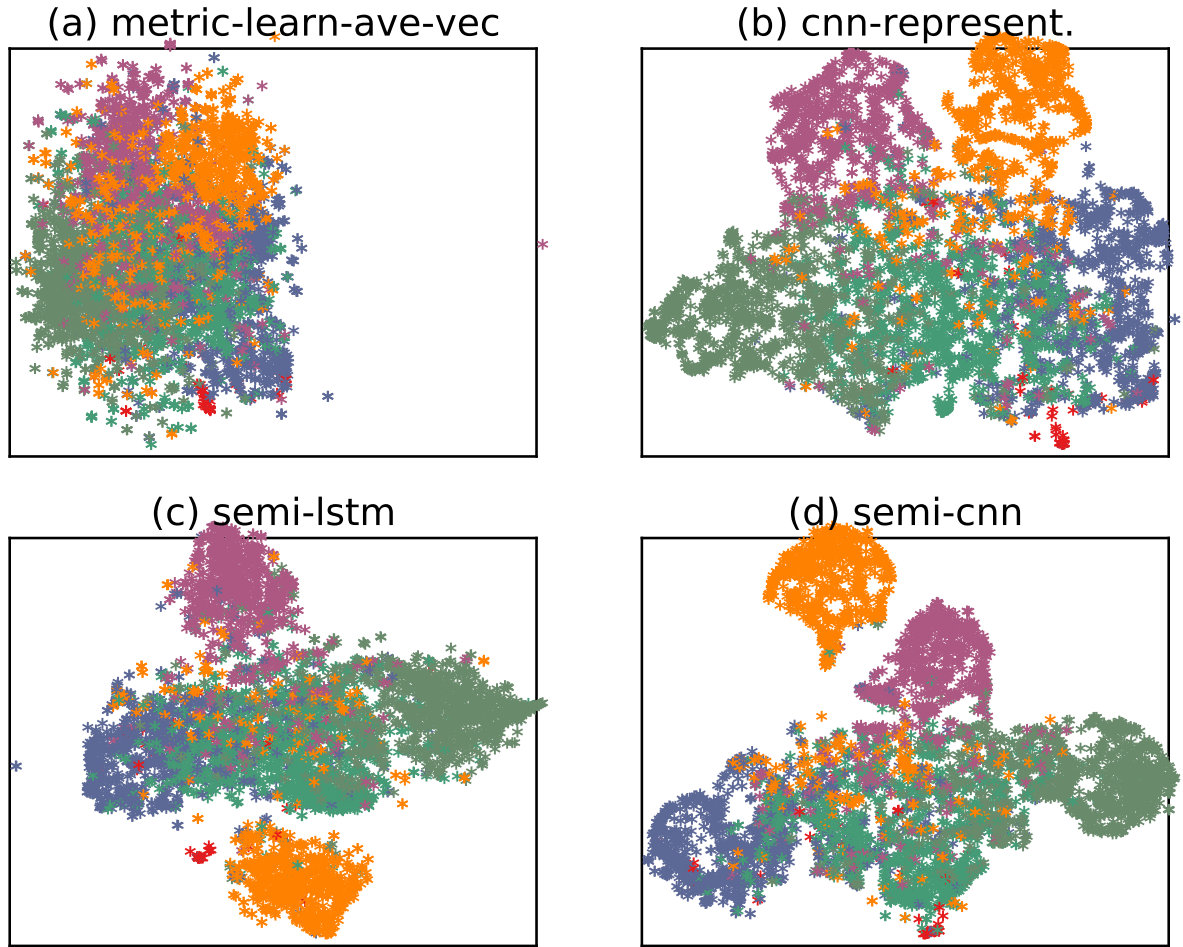


Figure 7: t-SNE visualizations of clustering results.

grouped all the systems into three categories: unsupervised (Unsup.), supervised (Sup.), and semi-supervised (Semisup.)<sup>3</sup>. We found that the supervised systems worked much better than the unsupervised counterparts, which implies that the small amount of labeled data is necessary for better performance. We also noticed that within the supervised systems, the systems using deep learning (CNN or LSTM) models worked better than the systems using metric learning method, which shows the power of deep learning models for short text modeling. Our “semi-cnn” system got the best performance on almost all the datasets.

Figure 7 visualizes clustering results on the *question\_type* dataset from four representative systems. In Figure 7(a), clusters severely overlap with each other. When using the CNN sentence representation model, we can clearly identify all clusters in Figure 7(b), but the boundaries between clusters are still obscure. The clustering results from our semi-supervised clustering algorithm are given in Figure 7(c) and Figure 7(d). We can see that the boundaries between clusters become much clearer. Therefore, our algorithm is very effective for short text clustering.

## 5 Related Work

Existing semi-supervised clustering methods fall into two categories: *constraint-based* and *representation-based*. In *constraint-based* methods (Davidson and Basu, 2007), some labeled information is used to constrain the clustering process. In *representation-based* methods (Bair, 2013), a representation model is first trained to satisfy the labeled information, and all data points are clustered based on representations from the representation model. Bilenko et al. (2004) proposed to integrate these two methods into a unified framework, which shares the same idea of our proposed method. However, they only employed the metric learning model for representation learning, which is a linear projection. Whereas, our method utilized deep learning models to learn representations in a more flexible non-linear space. Xu et al. (2015) also employed deep learning models for short text clustering. However, their method separated the representation learning

<sup>3</sup>All clustering systems are based on the same number of instances (total# in Table 3). For the semi-supervised and supervised systems, the labels for 1% of the instances are given (labeled# in Table 3). And the evaluation was conducted only on the unlabeled portion.

process from the clustering process, so it belongs to the *representation-based* method. Whereas, our method combined the representation learning process and the clustering process together, and utilized both labeled data and unlabeled data for representation learning and clustering.

## 6 Conclusion

In this paper, we proposed a semi-supervised clustering algorithm for short texts. We utilized deep learning models to learn representations for short texts, and employed a small amount of labeled data to specify our intention for clustering. We integrated the representation learning process and the clustering process into a unified framework, so that both of the two processes get some benefits from labeled data and unlabeled data. Experimental results on four datasets show that our method is more effective than other competitors.

## References

- Enrique Amigó, Julio Gonzalo, Javier Artilles, and Felisa Verdejo. 2009. A comparison of extrinsic clustering evaluation metrics based on formal constraints. *Information retrieval*, 12(4):461–486.
- David Arthur and Sergei Vassilvitskii. 2007. k-means++: The advantages of careful seeding. In *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 1027–1035.
- Eric Bair. 2013. Semi-supervised clustering methods. *Wiley Interdisciplinary Reviews: Computational Statistics*, 5(5):349–361.
- Somnath Banerjee, Krishnan Ramanathan, and Ajay Gupta. 2007. Clustering short texts using wikipedia. In *Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 787–788.
- Mikhail Bilenko, Sugato Basu, and Raymond J Mooney. 2004. Integrating constraints and metric learning in semi-supervised clustering. In *Proceedings of the twenty-first international conference on Machine learning*, page 11. ACM.
- Christopher M Bishop. 2006. *Pattern recognition and machine learning*. springer.
- Ian Davidson and Sugato Basu. 2007. A survey of clustering with instance level constraints. *ACM Transactions on Knowledge Discovery from Data*, 1:1–41.



- Inderjit S. Dhillon and Yuqiang Guan. 2003. Information theoretic clustering of sparse co-occurrence data. pages 517–520. IEEE Computer Society.
- Samah Fodeh, Bill Punch, and Pang-Ning Tan. 2011. On ontology-driven document clustering using core semantic features. *Knowledge and information systems*, 28(2):395–421.
- Alex Graves. 2012. *Supervised sequence labelling with recurrent neural networks*, volume 385.
- Geoffrey E Hinton and Ruslan R Salakhutdinov. 2006. Reducing the dimensionality of data with neural networks. *Science*, 313(5786):504–507.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, pages 1746–1751.
- Diederik Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Jens Lehmann, Robert Isele, Max Jakob, Anja Jentzsch, Dimitris Kontokostas, Pablo N Mendes, Sebastian Hellmann, Mohamed Morsey, Patrick van Kleef, Sören Auer, et al. 2014. Dbpedia-a large-scale, multilingual knowledge base extracted from wikipedia. *Semantic Web Journal*, 5:1–29.
- Xin Li and Dan Roth. 2002. Learning question classifiers. In *Proceedings of the 19th international conference on Computational linguistics-Volume 1*, pages 1–7.
- James MacQueen. 1967. Some methods for classification and analysis of multivariate observations. In *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*, volume 1, pages 281–297. Oakland, CA, USA.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- James Munkres. 1957. Algorithms for the assignment and transportation problems. *Journal of the Society for Industrial and Applied Mathematics*, 5(1):32–38.
- Nguyen Xuan Vinh, Julien Epps, and James Bailey. 2009. Information theoretic measures for clusterings comparison: is a correction for chance necessary? In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 1073–1080.
- Kilian Q Weinberger, John Blitzer, and Lawrence K Saul. 2005. Distance metric learning for large margin nearest neighbor classification. In *Advances in neural information processing systems*, pages 1473–1480.
- Jiaming Xu, Peng Wang, Guanhua Tian, Bo Xu, Jun Zhao, Fangyuan Wang, and Hongwei Hao. 2015. Short text clustering via convolutional neural networks. In *Proceedings of NAACL-HLT*, pages 62–69.
- Xiang Zhang and Yann LeCun. 2015. Text understanding from scratch. *arXiv preprint arXiv:1502.01710*.