

Hierarchical Heterogeneous Graph Representation Learning for Short Text Classification

Yaqing Wang¹, Song Wang^{1,2}, Quanming Yao^{*3}, Dejing Dou¹

¹Baidu Research, Baidu Inc., China

²Department of ECE, University of Virginia, USA

³Department of EE, Tsinghua University, China

{wangyaqing01, v_wangsong07, doudejing}@baidu.com

qyaoaa@tsinghua.edu.cn

Abstract

Short text classification is a fundamental task in natural language processing. It is hard due to the lack of context information and labeled data in practice. In this paper, we propose a new method called SHINE, which is based on graph neural network (GNN), for short text classification. First, we model the short text dataset as a hierarchical heterogeneous graph consisting of word-level component graphs which introduce more semantic and syntactic information. Then, we dynamically learn a short document graph that facilitates effective label propagation among similar short texts. Thus, comparing with existing GNN-based methods, SHINE can better exploit interactions between nodes of the same types and capture similarities between short texts. Extensive experiments on various benchmark short text datasets show that SHINE consistently outperforms state-of-the-art methods, especially with fewer labels.¹

1 Introduction

Short texts such as tweets, news feeds and web search snippets appear daily in our life (Pang and Lee, 2005; Phan et al., 2008). To understand these short texts, short text classification (STC) is a fundamental task which can be found in many applications such as sentiment analysis (Chen et al., 2019), news classification (Yao et al., 2019) and query intent classification (Wang et al., 2017).

STC is particularly hard in comparison to long text classification due to two key issues. The first key issue is that short texts only contain one or a few sentences whose overall length is small, which lack enough context information and strict syntactic structure to understand the meaning of texts (Tang et al., 2015; Wang et al., 2017). For example, it is hard to get the meaning of "Birthday girl is an amusing ride" without

knowing "Birthday girl" is a 2001 movie. A harder case is to understand a web search snippet such as "how much Tesla", which usually does not contain word order nor function words (Phan et al., 2008). In addition, real STC tasks usually only have a limited number of labeled data compared to the abundant unlabeled short texts emerging everyday (Hu et al., 2019). Therefore, auxiliary knowledge is required to understand short texts, examples include concepts that can be found in common sense knowledge graphs (Wang et al., 2017; Chen et al., 2019), latent topics extracted from the short text dataset (Hu et al., 2019), and entities residing in knowledge graphs (Hu et al., 2019). However, simply enriching auxiliary knowledge cannot solve the shortage of labeled data, which is another key issue commonly faced by real STC tasks (Pang and Lee, 2005; Phan et al., 2008). Yet the popularly used deep models require large-scale labeled data to train well (Kim, 2014; Liu et al., 2016).

Currently, graph neural networks (GNNs) designed for STC obtain the state-of-the-art performance (Hu et al., 2019; Ye et al., 2020). They both take the STC as the node classification problem on a graph with mixed nodes of different types: HGAT (Hu et al., 2019) builds a corpus-level graph modeling latent topics, entities and documents and STGCN (Ye et al., 2020) operates on a corpus-level graph of latent topics, documents and words. In both works, each document is connected to its nodes of a different type such as entities and latent topics but not to other documents. However, they do not fully exploit interactions between nodes of the same type. They also fail to capture the similarities between short documents, which is both useful to understand short texts (Zhu et al., 2003; Kenter and De Rijke, 2015; Wang et al., 2017) and an important to propagate few labels on graphs (Kipf and Welling, 2016). Besides, both works have large parameter sizes: HGAT (Hu

¹Codes are available at <https://github.com/tata1661/SHINE-EMNLP21>.

et al., 2019) is a GNN with dual-level attention and STGCN (Ye et al., 2020) merges the node representations with word embeddings obtained via a pretrained BERT (Devlin et al., 2019) via a bidirectional LSTM (Liu et al., 2016).

To address the aforementioned problems, we propose a novel Hierarchical heterogeneous graph representation learning method for STC called SHINE, which is able to fully exploit interactions between nodes of the same types and capture similarity between short texts. SHINE operates on a hierarchically organized heterogeneous corpus-level graph, which consists of the following graphs at different levels: (i) word-level component graphs model interactions between words, part-of-speech (POS) tags and entities which can be easily extracted and carry additional semantic and syntactic information to compensate for the lack of context information; and (ii) short document graph is dynamically learned and optimized to encode similarities between short documents which allows more effective label propagation among connected similar short documents. We conduct extensive experiments on a number of benchmark STC datasets including news, tweets, document titles and short reviews. Results show that the proposed SHINE consistently outperforms the state-of-the-art with a much smaller parameter size.

2 Related Works

2.1 Text Classification

Text classification assigns predefined labels to documents of variable lengths which may consist of a single or multiple sentences (Li et al., 2020). Traditional methods adopt a two-step strategy: first extract human-designed features such as bag-of-words (Blei et al., 2003) and term frequency-inverse document frequency (TF-IDF) (Aggarwal and Zhai, 2012) from documents, then learn classifiers such as support vector machine (SVM) (Cortes and Vapnik, 1995). Deep neural networks such as convolutional neural networks (CNN) (Kim, 2014) and long short-term memory (LSTM) (Liu et al., 2016) can directly obtain expressive representations from raw texts and conduct classification in an end-to-end manner.

Recently, graph neural networks (GNNs) (Defferrard et al., 2016; Kipf and Welling, 2016) have obtained the state-of-the-art performance on text classification. They can be divided into two types. The first type of GNNs constructs

document-level graphs where each document is modeled as a graph of word nodes, then formulates text classification as a whole graph classification problem (Defferrard et al., 2016). Examples are TLGNN (Huang et al., 2019), TextING (Zhang et al., 2020), HyperGAT (Ding et al., 2020), which establish word-word edges differently. In particular, some methods (Liu et al., 2019; Chen et al., 2020) propose to estimate the graph structure of the document-level graphs during learning. However, if only a few documents are labeled, these GNNs cannot work due to the lack of labeled graphs.

As is known, GNNs such as graph convolutional network (GCN) (Kipf and Welling, 2016) can conduct semi-supervised learning to solve node classification task on a graph where only a small number of nodes are labeled (Kipf and Welling, 2016). Therefore, another type of GNNs instead operates on a heterogeneous corpus-level graph which takes both text and word as nodes, and classifies unlabeled texts by node classification. Examples include TextGCN (Yao et al., 2019), TensorGCN (Liu et al., 2020), HeteGCN (Ragesh et al., 2021) and TG-Transformer (Zhang and Zhang, 2020) with different strategies to construct and handle heterogeneous nodes and edges. However, these methods cannot work well for short texts of limited length.

2.2 Short Text Classification (STC)

Short text classification (STC) is particularly challenging (Aggarwal and Zhai, 2012; Li et al., 2020). Due to limited length, short texts lack context information and strict syntactic structure which are vital to text understanding (Wang et al., 2017). Therefore, methods tailored for STC strive to incorporate various auxiliary information to enrich short text representations. Popularly used examples are concepts existing in external knowledge bases such as Probase (Wang et al., 2017; Chen et al., 2019) and latent topics discovered in the corpus (Zeng et al., 2018). However, simply enriching semantic information cannot compensate for the shortage of labeled data, which is a common problem faced by real short texts such as queries and online reviews (Pang and Lee, 2005; Phan et al., 2008). Thus, GNN-based methods which perform node classification for semi-supervised STC are utilized. HGAT (Hu et al., 2019) applies a GNN with dual-level attention to forward messages on a corpus-level graph modeling topics, entities

and documents jointly, where the entities are words linked to knowledge graphs. STGCN (Ye et al., 2020) operates on a corpus-level graph of topics, documents and words, and merges the node representations with word embeddings obtained via a pretrained BERT (Devlin et al., 2019) via a bidirectional LSTM (Liu et al., 2016). Currently, the state-of-the-art method on STC is HGAT (Hu et al., 2019; Yang et al., 2021).

3 Proposed Method

As mentioned in Section 2.2, GNN-based methods, i.e., HGAT and STGCN, can classify short texts while HGAT performs better. However, both works build a graph with mixed nodes of different types without fully exploiting interactions between nodes of the same type. Besides, they fail to capture the similarities between short documents, which can be important to propagate few labels on graphs. Here, we present the proposed SHINE which can address above limitations, thus is able to better compensate for the shortage of context information and labeled data for STC.

Given a short text dataset² \mathcal{S} containing short documents, we model \mathcal{S} as a hierarchically organized heterogeneous graph consisting of: (i) word-level component graphs: we construct word-level component graphs which model word-level semantic and syntactic information in order to compensate for the lack of context information; (ii) short document graph: we dynamically learn the short document graph via hierarchically pooling over word-level component graphs, such that the limited label information can be effectively propagated among similar short texts. A high-level illustration of SHINE is shown in Figure 1.

In the sequel, vectors are denoted by lowercase boldface, matrices by uppercase boldface. For a vector \mathbf{x} , $[\mathbf{x}]_i$ denotes the i th element of \mathbf{x} . For a matrix \mathbf{X} , \mathbf{x}^i denotes its i th row, $[\mathbf{X}]_{ij}$ denotes the (i, j) th entry of \mathbf{X} . For a set \mathcal{S} , $|\mathcal{S}|$ denotes the number of elements in \mathcal{S} .

3.1 Word-Level Component Graphs

To compensate for the lack of context information and syntactic structure in short documents, we leverage various word-level components which can bring in more syntactic and semantic information.

²For consistency, we refer each input text to classify as a document which may consist of one or a few sentences following (Tang et al., 2015; Hu et al., 2019).

Particularly, we consider the following three types of word-level components $\tau \in \{w, p, e\}$ in this paper: (i) word (w) which makes up short documents and carries semantic meaning; (ii) POS tag (p) which marks the syntactic role such as noun and verb of each word in the short text and is helpful for discriminating ambiguous words; and (iii) entity (e) which corresponds to word that can be found in auxiliary knowledge bases such that additional knowledge can be incorporated. SHINE can easily be extended with other components such as adding a topic graph on the first level. We use these three word-level components as they are well-known, easy to obtain at a low cost, and already surpass the state-of-the-art HGAT which use topics.

We first provide a general strategy to obtain node embeddings from different types of word-level component graphs, then describe in detail how to construct these graphs via common natural language processing techniques including tokenization, entity linking and POS tagging. In this way, we can fully exploit interactions between nodes of the same type.

3.1.1 Node Embedding Learning

Denote word-level component graph of type τ as $\mathcal{G}_\tau = \{\mathcal{V}_\tau, \mathbf{A}_\tau\}$ where \mathcal{V}_τ is a set of nodes and $\mathbf{A}_\tau \in \mathbb{R}^{|\mathcal{V}_\tau| \times |\mathcal{V}_\tau|}$ is the adjacency matrix. Each node $v_\tau^i \in \mathcal{V}_\tau$ is provided with node feature $\mathbf{x}_\tau^i \in \mathbb{R}^{d_\tau}$. For simplicity, the node features are collectively denoted as $\mathbf{X}_\tau \in \mathbb{R}^{|\mathcal{V}_\tau| \times d_\tau}$ with the i th row corresponds to one node feature \mathbf{x}_τ^i . These \mathcal{G}_τ s are used to capture the pairwise relationship between nodes of the same type, without being influenced by other types.

Provided with \mathcal{G}_τ and \mathbf{X}_τ , we use the classic 2-layer graph convolutional network (GCN) (Kipf and Welling, 2016) to obtain node embeddings \mathbf{H}_τ . Formally, \mathbf{H}_τ is updated as

$$\mathbf{H}_\tau = \tilde{\mathbf{A}}_\tau \cdot \text{ReLu}(\tilde{\mathbf{A}}_\tau \mathbf{X}_\tau \mathbf{W}_\tau^1) \mathbf{W}_\tau^2, \quad (1)$$

where $[\text{ReLu}(\mathbf{x})]_i = \max([\mathbf{x}]_i, 0)$, $\tilde{\mathbf{A}}_\tau = \mathbf{D}_\tau^{-\frac{1}{2}} (\mathbf{I} + \mathbf{A}_\tau) \mathbf{D}_\tau^{-\frac{1}{2}}$ with $[\mathbf{D}_\tau]_{ii} = \sum_j [\mathbf{A}_\tau]_{ij}$, and $\mathbf{W}_\tau^1, \mathbf{W}_\tau^2$ are trainable parameters.

3.1.2 Graph Construction

Next, we present the details of how to construct each \mathcal{G}_τ from \mathcal{S} .

Word Graph \mathcal{G}_w . We construct a word graph $\mathcal{G}_w = \{\mathcal{V}_w, \mathbf{A}_w\}$ where word nodes are connected based on local co-occurrence relationships, while

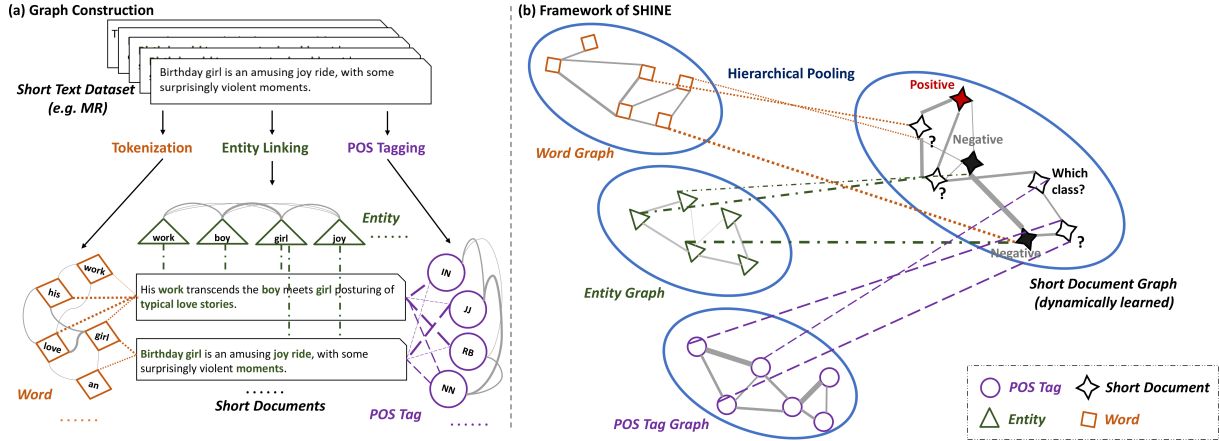


Figure 1: A high-level illustration of (a) how we construct a heterogeneous corpus-level graph from a short text dataset using well-known natural language processing techniques; and (b) framework of the proposed SHINE which hierarchically pools over word-level component graphs to obtain short document graph where node classification is conducted to classify those unlabeled nodes. SHINE is trained end-to-end on the complete two-level graph with respect to the classification loss. The plotted examples of short texts are taken from the movie review (MR) dataset (Pang and Lee, 2005).

other types of relationship such as syntactic dependency (Liu et al., 2020) can also be used. We set $[A_w]_{ij} = \max(\text{PMI}(v_w^i, v_w^j), 0)$ where PMI denotes the point-wise mutual information between words $v_w^i, v_w^j \in \mathcal{V}_w$ (Yao et al., 2019). We initialize the node feature $\mathbf{x}_w^i \in \mathbb{R}^{|\mathcal{V}_w|}$ for $v_w^i \in \mathcal{V}_w$ as a one-hot vector. Once learned by (1), \mathbf{H}_w is able to encode the topological structure of \mathcal{G}_w which is specific to \mathcal{S} . We can also leverage generic semantic information by concatenating \mathbf{H}_w with pretrained word embeddings $\hat{\mathbf{H}}_w$ extracted from large text corpus such as Wikidata (Vrandečić and Krötzsch, 2014).

POS Tag Graph \mathcal{G}_p . We use the default POS tag set of NLTK³ to obtain the POS tag for each word of short text in \mathcal{S} , which forms the POS tag node set \mathcal{V}_p . Similar to \mathcal{G}_w , we construct a co-occurrence POS tag graph $\mathcal{G}_p = \{\mathcal{V}_p, \mathbf{A}_p\}$ with $[A_p]_{ij} = \max(\text{PMI}(v_p^i, v_p^j), 0)$, where the inputs are POS tags for all words. Then we again initialize the node feature $\mathbf{x}_p^i \in \mathbb{R}^{|\mathcal{V}_p|}$ as a one-hot vector.

Entity Graph \mathcal{G}_e . We obtain the entity node set \mathcal{V}_e by recognizing entities presented in the NELL knowledge base (Carlson et al., 2010). In contrast to words and POS tags which are abundant in the documents, the number of entities is much smaller. Most short documents only contain one entity, which makes it infeasible to calculate co-occurrence statistics between entities. Instead, we

first learn the entity feature $\mathbf{x}_e^i \in \mathbb{R}^{d_e}$ of each $v_e^i \in \mathcal{G}_e$ from NELL, using the classic knowledge graph embedding method TransE (Bordes et al., 2013). Then, we measure the cosine similarity $c(v_e^i, v_e^j)$ between each entity pair $v_e^i, v_e^j \in \mathcal{V}_e$ and set $[A_e]_{ij} = \max(c(x_e^i, x_e^j), 0)$.

3.2 Short Document Graph

As discussed in Section 2.1, the reason why GNN-based methods, which take short documents classification as node classification tasks, can deal with few labels is the usage of adjacent matrix which models the similarities between short documents. However, STGCN and HGAT do not consider such similarities.

Here, to effectively propagate the limited label information, we dynamically learn the short document graph $\mathcal{G}_s = \{\mathcal{V}_s, \mathbf{A}_s\}$ based on embeddings pooled over word-level component graph to encode the similarity between short documents, where $v_s^i \in \mathcal{V}_s$ corresponds to one short document in \mathcal{S} , and \mathbf{A}_s is the learned adjacency matrix. As shown in Figure 1, we propose to obtain \mathcal{G}_s via hierarchically pooling over word-level component graphs, hence \mathcal{G}_s is dynamically learned and optimized during training. This learned \mathcal{G}_s then facilitates efficient label propagation among connected short texts.

3.2.1 Hierarchical Pooling over \mathcal{G}_τ s

In this section, we propose to learn \mathbf{A}_s via a text-specific hierarchically pooling over multiple word-

³<http://www.nltk.org>

level component graphs (\mathcal{G}_τ s).

With \mathbf{H}_τ obtained from (1), we represent each $v_s^i \in \mathcal{G}_s$ by pooling over node embeddings in \mathcal{G}_τ with respect to \mathcal{G}_τ :

$$\hat{\mathbf{x}}_\tau^i = u(\mathbf{H}_\tau^\top \mathbf{s}_\tau^i), \quad (2)$$

where superscript $(\cdot)^\top$ denotes the transpose operation, and $u(\mathbf{x}) = \mathbf{x}/\|\mathbf{x}\|_2$ normalizes \mathbf{x} to unit norm. Particularly, each \mathbf{s}_τ^i is computed as follows:

- When $\tau = w$ or p : $[\mathbf{s}_\tau^i]_j = \text{TF-IDF}(v_\tau^j, v_s^i)$ where TF-IDF is the term frequency-inverse document frequency (Aggarwal and Zhai, 2012). We then normalize \mathbf{s}_τ^i as $\mathbf{s}_\tau^i / \sum_{j=1} [\mathbf{s}_\tau^i]_j$.
- When $\tau = e$: $[\mathbf{s}_\tau^i]_j = 1$ if v_e^j exists in v_s^i and 0 otherwise, as most short texts only contain one entity. \mathbf{s}_τ^i is then normalized as $\mathbf{s}_\tau^i / \sum_{j=1} [\mathbf{s}_\tau^i]_j$.

This can be seen as explaining each short document from the perspective from words, POS tags and entities collectively. Finally, we obtain the short text representation \mathbf{x}_s^i of v_s^i is obtained as

$$\mathbf{x}_s^i = \hat{\mathbf{x}}_w^i \parallel \hat{\mathbf{x}}_p^i \parallel \hat{\mathbf{x}}_e^i, \quad (3)$$

where \parallel means concatenating vectors along the last dimension. Please note that concatenation is just an instantiation which already obtains good performance. It can be replaced by more complex aggregation function such as weighted average or LSTM.

3.2.2 Dynamic Graph Learning

Now, we obtain \mathbf{A}_s on the fly using the learned short document features \mathbf{x}_s^i 's:

$$[\mathbf{A}_s]_{ij} = \begin{cases} (\mathbf{x}_s^i)^\top \mathbf{x}_s^j & \text{if } (\mathbf{x}_s^i)^\top \mathbf{x}_s^j \geq \delta_s, \\ 0 & \text{otherwise} \end{cases}, \quad (4)$$

where δ_s is a threshold used to sparsity \mathbf{A}_s such that short documents are connected only if they are similar enough viewed from the perspective of \mathcal{G}_τ s. Note that the resultant \mathcal{G}_s is dynamically changing along with the optimization process, where \mathbf{H}_τ , \mathbf{x}_s^i , and \mathbf{A}_s are all optimized and improved.

Upon this \mathcal{G}_s , we propagate label information among similar short documents via a 2-layer GCN. Let \mathbf{X}_s collectively record short text embeddings with \mathbf{x}_s^i on the i th row. The class predictions of all short documents in \mathcal{S} with respect to C classes are obtained as

$$\hat{\mathbf{Y}}_s = \text{softmax}(\mathbf{A}_s \cdot \text{ReLu}(\mathbf{A}_s \mathbf{X}_s \mathbf{W}_s^1) \cdot \mathbf{W}_s^2), \quad (5)$$

where $[\text{softmax}(\mathbf{x})]_i = \exp([\mathbf{x}]_i) / \sum_j \exp([\mathbf{x}]_j)$ is applied for each row, \mathbf{W}_s^1 and \mathbf{W}_s^2 are trainable parameters.

We train the complete model by optimizing the cross-entropy loss function in an end-to-end manner:

$$\mathcal{L} = - \sum_{i \in \mathcal{I}_l} (\mathbf{y}_s^i)^\top \log(\hat{\mathbf{y}}_s^i), \quad (6)$$

where \mathcal{I}_l records the indices of the labeled short documents, $\mathbf{y}_s^i \in \mathbb{R}^C$ is a one-hot vector with all 0s but a single one denoting the index of the ground truth class $c \in \{1, \dots, C\}$. By jointly optimized with respect to the single objective, different types of graphs can influence each other. During learning, node embeddings of \mathcal{G}_τ for all $\tau \in \{w, p, e, s\}$ and \mathbf{A}_s are all updated. The complete procedure of SHINE is shown in Algorithm 1.

Algorithm 1 SHINE Algorithm.

Input: short text dataset \mathcal{S} , word-level component graphs $\mathcal{G}_\tau = \{\mathcal{V}_\tau, \mathbf{A}_\tau\}$ with node features \mathbf{X}_τ , sample-specific aggregation vectors $\{\mathbf{s}_\tau^i\}$ where $\tau \in \{w, p, e\}$;

- 1: **for** $t = 1, 2, \dots, T$ **do**
- 2: **for** $\tau \in \{w, p, e\}$ **do**
- 3: obtain node embeddings \mathbf{H}_τ of \mathcal{G}_τ by (1);
- 4: **end for**
- 5: obtain short document features \mathbf{X}_s via hierarchically pooling over \mathcal{G}_τ s by (3);
- 6: obtain short document embeddings from \mathcal{G}_s and make the class prediction by (5);
- 7: optimize model parameter with respect to (6) by back propagation;
- 8: **end for**

4 Experiments

All results are averaged over five runs and are obtained on a PC with 32GB memory, Intel-i8 CPU, and a 32GB NVIDIA Tesla V100 GPU.

4.1 Datasets

We perform experiments on a variety of publicly accessible benchmark short text datasets (Table 1):

- (i) **Ohsumed**⁴: a subset of the bibliographic Ohsumed dataset (Hersh et al., 1994) used in (Hu et al., 2019) where the title is taken as the short text to classify.
- (ii) **Twitter**⁵: a collection of tweets expressing

⁴<http://disi.unitn.it/moschitti/corpora.htm>

⁵http://www.nltk.org/howto/twitter.html#corpus_reader

	# texts	avg. length	# classes	# train (ratio)	# words	# entities	# POS tags
Ohsumed	7,400	6.8	23	460 (6.22%)	11,764	4,507	38
Twitter	10,000	3.5	2	40 (0.40%)	21,065	5,837	41
MR	10,662	7.6	2	40 (0.38%)	18,764	6,415	41
Snippets	12,340	14.5	8	160 (1.30%)	29040	9737	34
TagMyNews	32,549	5.1	7	140 (0.43%)	38629	14734	42

Table 1: Summary of short text datasets used.

positive or negative attitude towards some contents.

- (iii) **MR**⁶: a movie review dataset for sentiment analysis (Pang and Lee, 2005).
- (iv) **Snippets**⁷: a dataset of web search snippets returned by Google Search (Phan et al., 2008).
- (v) **TagMyNews**: a dataset contains English news titles collected from Really Simple Syndication (RSS) feeds, as adopted by Hu et al. (2019).

We tokenize each sentence and remove stopping words and low-frequency words which appear less than five times in the corpus as suggested in (Yao et al., 2019; Hu et al., 2019).

Following (Hu et al., 2019), we randomly sample 40 labeled short documents from each class where half of them forms the training set and the other half forms the validation set for hyperparameter tuning. The rest short documents are taken as the test set, which are unlabeled during training.

4.2 Compared Methods

The proposed **SHINE** is compared with the following methods.

- *Group (A)*. Two-step feature extraction and classification methods include (i) **TF-IDF+SVM** and (ii) **LDA+SVM** (Cortes and Vapnik, 1995) which use support vector machine to classify documents represented by TF-IDF feature and LDA feature respectively; and (iii) **PTE**⁸ (Tang et al., 2015) which learns a linear classifier upon documents represented as the average word embeddings pretrained from bipartite word-word, word-document and word-label graphs.

⁶<http://www.cs.cornell.edu/people/pabo/movie-review-data/>

⁷Snippets and TagMyNews are downloaded from <http://acube.di.unipi.it:80/tmn-dataset/>.

⁸<https://github.com/mnqu/PTE>

- *Group (B)*. **BERT**⁹ (Devlin et al., 2019) which is pretrained on a large corpus and fine-tuned together with a linear classifier for the short text classification task. Each document is represented as the averaged word embeddings (denote as **-avg**) or the embedding of the CLS token (denote as **-CLS**).
- *Group (C)*. Deep text classification methods include (i) **CNN** (Kim, 2014) and (ii) **LSTM** (Liu et al., 2016) where the input word embeddings are either randomly initialized (denote as **-rand**) or pretrained from large text corpus (denote as **-pre**); GNNs which perform graph classification on document-level graphs including (iii) **TL-GNN**¹⁰ (Huang et al., 2019), (iv) **TextING**¹¹ (Zhang et al., 2020), and (v) **HyperGAT**¹² (Ding et al., 2020); GNNs which perform node classification on corpus-level graphs including (vi) **TextGCN**¹³ (Yao et al., 2019) and (vii) **TensorGCN**¹⁴ (Liu et al., 2020). HeteGCN (Ragesh et al., 2021) and TG-Transformer (Zhang and Zhang, 2020) are not compared due to the lack of publicly available codes.
- *Group (D)*. Deep STC methods including (i) **STCKA**¹⁵ (Chen et al., 2019): an attention-based BiLSTM model, which fuses concept found in auxiliary knowledge bases into short document embedding; (ii) **HGAT**¹⁶ (Hu et al., 2019) which operates on a corpus-level graph of entities, topics and documents using a GNN with dual-level attention; and (iii) **STGCN**¹⁷ (Ye et al., 2020) which operates on a corpus-level graph of words, topics and documents and uses a bidirectional LSTM to merge the word

⁹https://tfhub.dev/tensorflow/bert_en_uncased_L-12_H-768_A-12/4

¹⁰<https://github.com/LindgeW/TextLevelGNN>

¹¹<https://github.com/CRIPAC-DIG/TextING>

¹²<https://github.com/kaize0409/HyperGAT>

¹³https://github.com/yao8839836/text_gcnn

¹⁴<https://github.com/THUMLP/TensorGCN>

¹⁵<https://github.com/AIRobotZhang/STCKA>

¹⁶<https://github.com/ytc272098215/HGAT>

¹⁷<https://github.com/yzhiahao/STGCN>

Group	Model	Ohsumed		Twitter		MR		Snippets		TagMyNews	
		ACC	F1	ACC	F1	ACC	F1	ACC	F1	ACC	F1
(A)	TFIDF+SVM	39.02	24.78	53.69	52.45	54.29	48.13	64.70	59.17	39.91	32.05
	LDA+SVM	38.61	25.03	54.34	53.97	54.40	48.39	62.54	56.40	40.40	30.40
	PTE	36.63	19.24	54.24	53.17	54.74	52.36	63.10	58.96	40.32	33.56
(B)	BERT-avg	23.91	4.98	54.92	51.15	51.69	50.65	79.31	78.47	55.13	44.26
	BERT-CLS	21.76	4.81	52.00	43.34	53.48	46.99	81.53	<u>79.03</u>	58.17	41.04
(C)	CNN-rand	35.25	13.95	52.58	51.91	54.85	51.23	48.34	42.12	28.76	15.82
	CNN-pre	32.92	12.06	56.34	55.86	58.32	57.99	77.09	69.28	57.12	45.37
	LSTM-rand	23.30	5.20	54.81	53.85	53.13	52.98	30.74	25.04	25.89	17.01
	LSTM-pre	29.05	5.09	58.20	58.16	59.73	59.19	75.07	67.31	53.96	42.14
(D)	TLGNN	35.76	13.12	58.33	53.86	58.48	58.45	70.25	63.18	44.43	32.33
	TextING	38.27	21.34	59.79	59.44	58.89	58.76	71.13	70.71	52.53	40.20
	HyperGAT	36.60	19.98	58.42	53.71	58.65	58.62	70.89	63.42	45.60	31.51
	TextGCN	41.56	<u>27.43</u>	60.15	59.82	59.12	58.98	77.82	71.95	54.28	46.01
	TensorGCN	41.84	24.24	61.24	61.19	59.22	58.78	74.38	73.96	55.58	43.21
	STCKA	30.19	10.12	57.45	56.97	53.22	50.11	68.96	61.27	30.44	20.01
	HGAT	<u>42.68</u>	24.82	63.21	62.48	<u>62.75</u>	<u>62.36</u>	<u>82.36</u>	74.44	<u>61.72</u>	<u>53.81</u>
	STGCN	33.91	27.22	<u>64.33</u>	<u>64.29</u>	58.18	58.11	70.01	69.93	34.74	34.01
	SHINE (ours)	45.57	30.98	72.54	72.19	64.58	63.89	82.39	81.62	62.50	56.21
	relative \uparrow (%)	6.77	12.94	12.76	12.29	2.92	2.45	0.85	3.17	1.26	4.46

Table 2: Test performance (%) measured on short text datasets. The best results (according to the pairwise t-test with 95% confidence) are highlighted in bold. The second best results are marked in italic. The last row records the relative improvement (%) of SHINE over the second best result.

embeddings learned by a GNN with word embeddings produced by a pretrained BERT.

For these baseline methods, we either show the results reported in previous research (Hu et al., 2019; Yang et al., 2021) or run the public codes provided by the authors. For fairness, we use the public 300-dimensional GloVe word embeddings¹⁸ in all methods which require pretrained word embeddings (Pennington et al., 2014).

Hyperparameter Setting. For all methods, we find hyperparameters using the validation set via grid search. For SHINE, we set entity embedding dimension d_e as 100. For all the datasets, we set the sliding window size of PMI as 5 for both \mathcal{G}_w and \mathcal{G}_p , set the embedding size of all GCN layers used in SHINE as 200, and set the threshold δ_s for \mathcal{G}_s as 2.5. We implement SHINE in PyTorch and train the model for a maximum number of 1000 epochs using Adam (Kingma and Ba, 2014) with learning rate 10^{-3} . We early stop training if the validation loss does not decrease for 10 consecutive epochs. Dropout rate is set as 0.5.

¹⁸<http://nlp.stanford.edu/data/glove.6B.zip>

Evaluation Metrics. We evaluate the classification performance using test accuracy (denote as ACC in short) and macro-averaged F1 score (denote as F1 in short) following (Tang et al., 2015; Yang et al., 2021).

4.3 Benchmark Comparison

Performance Comparison. Table 2 shows the performance. As can be seen, GNN-based methods in group (D) obtain better classification results in general, where the proposed SHINE consistently obtains the state-of-the-art test accuracy and macro-F1 score. This can be attributed to the effective semantic and syntactic information fusion and the modeling of short document graphs.

In addition, if we order datasets by increasing average text length (i.e., Twitter, TagMyNews, Ohsumed, MR and Snippets), we can find that SHINE basically obtains larger relative improvement over the second best method on shorter documents as shown in the last row of Table 2. This validates the efficacy of label propagation in SHINE, which can be attributed to the dynamical learning of short document graph. As shown,

GNNs which perform node classification on the corpus-level graph obtain better performance than GNNs which perform graph classification on short text datasets with a few labeled short documents. Another common observation is that incorporating pretrained word embeddings can consistently improve the accuracy, as can be observed by comparing CNN-pre to CNN-rand, LSTM-pre to LSTM-rand, BiLSTM-pre to BiLSTM-rand. CNN and LSTM can obtain worse performance than traditional methods in group (A), such as results on Ohsumed. The fine-tuned BERT encodes generic semantic information from a large corpus, but it cannot beat SHINE which is particularly designed to handle the short text dataset.

Model Size Comparison. Table 3 presents the parameter size of SHINE and the two most relevant GNN-based methods, i.e., HGAT and STGCN. As can be seen, SHINE takes much smaller parameter size. The reason is that instead of organizing different types of nodes in the same graph like HGAT and STGCN, SHINE separately constructs graphs for each type of nodes and pools from them to represent short documents. Thus, the graph used in SHINE can be much smaller than HGAT and STGCN, which leads to a reduction of the parameter number. We also observe that SHINE takes less training time per epoch.

	HGAT	STGCN	SHINE
Ohsumed	3,091,523	2,717,104	212,146
Twitter	2,312,672	3,201,824	201,604
MR	6,030,640	3,326,224	201,604
Snippets	8,892,778	3,238,304	204,616
TagMyNews	10,162,899	6,653,024	204,114

Table 3: Comparison of model parameter size.

4.4 Ablation Study

We compare with different variants of SHINE to evaluate the contribution of each part:

- (i) **w/o \mathcal{G}_w , w/o \mathcal{G}_p and w/o \mathcal{G}_e :** remove one single \mathcal{G}_τ from SHINE while keeping the other parts unchanged.
- (ii) **w/o pre:** do not concatenate \mathbf{H}_w with pretrained word embeddings $\hat{\mathbf{H}}_w$.
- (iii) **w/ pre \mathbf{X}_w :** initializes node embeddings \mathbf{X}_w of \mathcal{G}_w as pretrained word embeddings $\hat{\mathbf{H}}_w$ directly.

- (iv) **w/o word GNN:** fix \mathbf{H}_τ as the input node features \mathbf{X}_τ of each cGH_τ , therefore the node embeddings of \mathcal{G}_s are simply weighted average of corresponding word-level features.
- (v) **w/o doc GNN:** use label propagation (Zhou et al., 2004) to directly obtain class prediction using \mathbf{A}_s learned by (4) and \mathbf{x}_s^i learned by (3).
- (vi) **w/ a single GNN:** run a single GNN on a heterogeneous corpus-level graph containing the same set of words, entities, POS tags and documents as ours. We modify TextGCN (Yao et al., 2019) to handle this case.

Table 4 shows the results. As shown, word-level component graphs and short document graph contribute larger to the effectiveness of SHINE. The concatenation of pretrained word embedding can slightly improve the performance. However, “w/ pre \mathbf{X}_w ” is worse than SHINE. This shows the benefits of separating corpus-specific and general semantic information: using \mathcal{G}_w with one-hot initialized features to capture corpus-specific topology among words, while using pretrained word embeddings to bring in general semantic information extracted from an external large corpus. The performance gain of SHINE with respect to “w/o word GNN” validates the necessity of (i) message passing among nodes of the same type and update node embeddings accordingly and (ii) update \mathcal{G}_τ s with respect to the STC task. While the improvement of SHINE upon “w/o doc GNN” shows that refining short document embeddings by GNN is useful. Finally, SHINE defeats “w/ a single GNN” which uses the same amount of information. This reveals that SHINE outperforms due to model design. Figure 2 further plots the influences of incrementally adding in more word-level component graphs and the short document graph. This again validates the effectiveness of SHINE framework.




Figure 2: Effect of adding each graph component.

Model	Ohsumed		Twitter		MR		Snippets		TagMyNews	
	ACC	F1	ACC	F1	ACC	F1	ACC	F1	ACC	F1
w/o \mathcal{G}_w	21.91	11.87	60.93	60.39	55.03	54.00	70.25	68.82	55.68	48.38
w/o \mathcal{G}_p	26.89	13.10	67.37	66.78	60.21	59.79	77.66	75.86	60.37	52.96
w/o \mathcal{G}_e	30.17	15.31	68.46	67.89	61.54	60.60	80.48	77.82	60.44	54.10
w/o \mathcal{G}_s	33.20	18.93	68.53	68.19	61.08	60.87	78.68	77.74	61.05	54.51
w/o pre	36.23	21.50	69.04	68.57	61.97	61.31	78.47	78.01	61.09	54.57
w/ pre \mathbf{X}_w	27.13	19.94	66.70	66.35	60.56	60.51	71.70	70.52	57.04	50.11
w/o word GNN	25.60	17.14	54.82	53.85	53.49	52.74	65.43	64.62	54.96	45.32
w/o doc GNN	37.41	25.85	70.60	70.47	61.94	61.60	79.27	78.10	61.38	55.93
w/ a single GNN	42.56	28.18	61.35	61.20	60.39	60.21	78.52	73.64	56.58	48.18
SHINE	45.57	30.98	72.54	72.19	64.58	63.89	82.39	81.62	62.50	56.21

Table 4: Test performance (%) of SHINE and its variants on short text datasets. The best results (according to the pairwise t-test with 95% confidence) are highlighted in bold.




Figure 3: Model sensitivity analysis of SHINE on Twitter.

4.5 Model Sensitivity

We further examine the impact of labeled training data proportion for GNN-based methods which perform node classification on the corpus-level graph, including TextGCN, TensorGCN, HGAT, STGCN and SHINE. Figure 3(a) plots the results. As shown, SHINE consistently outperforms other methods, where the performance gap is increasing with fewer labeled training data. Figure 3(b) plots the impact of threshold δ_s in (4). At first, performance increases with a larger δ_s which leads to a sparse \mathcal{G}_s where only certainly similar short documents are connected to propagate information. However, when δ_s is too large, \mathcal{G}_s loses its functionality and reduces to w/o \mathcal{G}_τ in Table 4. Finally, recall that we set the embedding size of all GCN layers used in SHINE equally. Figure 3(c) plots the effect of varying this embedding size. As observed, small embedding size cannot capture enough information while a overly large embedding size may not improve the performance but is more computational costly.

5 Conclusion

In this paper, we propose SHINE, a novel hierarchical heterogeneous graph representation learning method for short text classification. It is particularly useful to compensate for the lack of context information and propagate the limited number of labels efficiently. Specially, SHINE can effectively learn from a hierarchical graph modeling different perspectives of the short text dataset: word-level component graphs are used to understand short texts from the semantic and syntactic perspectives, and the dynamically learned short document graph allows efficient and effective label propagation among similar short documents. Extensive experiments show that SHINE outperforms the others consistently.

Acknowledgements

We thank the anonymous reviewers for their valuable comments. Parts of experiments were carried out on Baidu Data Federation Platform. Correspondence author is Quanming Yao.

References

- Charu C Aggarwal and ChengXiang Zhai. 2012. A survey of text classification algorithms. In *Mining Text Data*, pages 163–222. Springer.
- David M Blei, Andrew Y Ng, and Michael I Jordan. 2003. Latent dirichlet allocation. *Journal of Machine Learning Research*, 3(Jan):993–1022.
- Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. Translating embeddings for modeling multi-relational data. In *Advances in Neural Information Processing Systems*, pages 2787–2795.
- Andrew Carlson, Justin Betteridge, Bryan Kisiel, Burr Settles, Estevam R Hruschka Jr, and Tom M Mitchell. 2010. Toward an architecture for never-ending language learning. In *AAAI Conference on Artificial Intelligence*, pages 1306–1313.
- Jindong Chen, Yizhou Hu, Jingping Liu, Yanghua Xiao, and Haiyun Jiang. 2019. Deep short text classification with knowledge powered attention. In *AAAI Conference on Artificial Intelligence*, pages 6252–6259.
- Yu Chen, Lingfei Wu, and Mohammed Zaki. 2020. Iterative deep graph learning for graph neural networks: Better and robust node embeddings. In *Advances in Neural Information Processing Systems*, pages 19314–19326.
- Corinna Cortes and Vladimir Vapnik. 1995. Support-vector networks. *Machine Learning*, 20(3):273–297.
- Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. 2016. Convolutional neural networks on graphs with fast localized spectral filtering. In *Advances in Neural Information Processing Systems*, pages 3844–3852.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Kaize Ding, Jianling Wang, Jundong Li, Dingcheng Li, and Huan Liu. 2020. [Be more with less: Hypergraph attention networks for inductive text classification](#). In *Conference on Empirical Methods in Natural Language Processing*, pages 4927–4936, Online. Association for Computational Linguistics.
- William Hersh, Chris Buckley, TJ Leone, and David Hickam. 1994. Ohsumed: An interactive retrieval evaluation and new large test collection for research. In *International ACM-SIGIR Conference on Research and Development in Information Retrieval*, pages 192–201.
- Linmei Hu, Tianchi Yang, Chuan Shi, Houye Ji, and Xiaoli Li. 2019. [Heterogeneous graph attention networks for semi-supervised short text classification](#). In *Conference on Empirical Methods in Natural Language Processing*, pages 4821–4830, Hong Kong, China. Association for Computational Linguistics.
- Lianzhe Huang, Dehong Ma, Sujian Li, Xiaodong Zhang, and Houfeng Wang. 2019. [Text level graph neural network for text classification](#). In *Conference on Empirical Methods in Natural Language Processing*, pages 3444–3450, Hong Kong, China. Association for Computational Linguistics.
- Tom Kenter and Maarten De Rijke. 2015. Short text similarity with word embeddings. In *ACM international on conference on information and knowledge management*, pages 1411–1420.
- Yoon Kim. 2014. [Convolutional neural networks for sentence classification](#). In *Conference on Empirical Methods in Natural Language Processing*, pages 1746–1751, Doha, Qatar. Association for Computational Linguistics.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*.
- Thomas N Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*.
- Qian Li, Hao Peng, Jianxin Li, Congyin Xia, Renyu Yang, Lichao Sun, Philip S Yu, and Lifang He. 2020. A survey on text classification: From shallow to deep learning. *arXiv preprint arXiv:2008.00364*.
- Pengfei Liu, Shuaichen Chang, Xuanjing Huang, Jian Tang, and Jackie Chi Kit Cheung. 2019. Contextualized non-local neural networks for sequence learning. In *AAAI Conference on Artificial Intelligence*, pages 6762–6769.
- Pengfei Liu, Xipeng Qiu, and Xuanjing Huang. 2016. Recurrent neural network for text classification with multi-task learning. In *International Joint Conference on Artificial Intelligence*, pages 2873–2879.
- Xien Liu, Xinxin You, Xiao Zhang, Ji Wu, and Ping Lv. 2020. Tensor graph convolutional networks for text classification. In *AAAI Conference on Artificial Intelligence*, pages 8409–8416.
- Bo Pang and Lillian Lee. 2005. [Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales](#). In *Annual Meeting of the Association for Computational Linguistics*, pages 115–124, Ann Arbor, Michigan. Association for Computational Linguistics.

- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. [GloVe: Global vectors for word representation](#). In *Conference on Empirical Methods in Natural Language Processing*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- Xuan-Hieu Phan, Le-Minh Nguyen, and Susumu Horiguchi. 2008. Learning to classify short and sparse text & web with hidden topics from large-scale data collections. In *International Conference on World Wide Web*, pages 91–100.
- Rahul Ragesh, Sundararajan Sellamanickam, Arun Iyer, Ramakrishna Bairi, and Vijay Lingam. 2021. HeteGCN: Heterogeneous graph convolutional networks for text classification. In *ACM International Conference on Web Search and Data Mining*, pages 860–868.
- Jian Tang, Meng Qu, and Qiaozhu Mei. 2015. PTE: Predictive text embedding through large-scale heterogeneous text networks. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1165–1174.
- D. Vrandečić and M. Krötzsch. 2014. Wikidata: A free collaborative knowledgebase. *Communications of the ACM*, 57(10):78–85.
- Jin Wang, Zhongyuan Wang, Dawei Zhang, and Jun Yan. 2017. Combining knowledge with deep convolutional neural networks for short text classification. In *International Joint Conference on Artificial Intelligence*, pages 2915–2921.
- Tianchi Yang, Linmei Hu, Chuan Shi, Houye Ji, Xiaoli Li, and Liqiang Nie. 2021. HGAT: Heterogeneous graph attention networks for semi-supervised short text classification. *ACM Transactions on Information Systems*, 39(3):1–29.
- Liang Yao, Chengsheng Mao, and Yuan Luo. 2019. Graph convolutional networks for text classification. In *AAAI Conference on Artificial Intelligence*, volume 33, pages 7370–7377.
- Zhihao Ye, Gongyao Jiang, Ye Liu, Zhiyong Li, and Jin Yuan. 2020. Document and word representations generated by graph convolutional network and bert for short text classification. In *European Conference on Artificial Intelligence*, pages 2275–2281.
- Jichuan Zeng, Jing Li, Yan Song, Cuiyun Gao, Michael R. Lyu, and Irwin King. 2018. [Topic memory networks for short text classification](#). In *Conference on Empirical Methods in Natural Language Processing*, pages 3120–3131, Brussels, Belgium. Association for Computational Linguistics.
- Haopeng Zhang and Jiawei Zhang. 2020. [Text graph transformer for document classification](#). In *Conference on Empirical Methods in Natural Language Processing*, pages 8322–8327, Online. Association for Computational Linguistics.
- Yufeng Zhang, Xueli Yu, Zeyu Cui, Shu Wu, Zhongzhen Wen, and Liang Wang. 2020. [Every document owns its structure: Inductive text classification via graph neural networks](#). In *Annual Meeting of the Association for Computational Linguistics*, pages 334–339, Online. Association for Computational Linguistics.
- Dengyong Zhou, Olivier Bousquet, Thomas N Lal, Jason Weston, and Bernhard Schölkopf. 2004. Learning with local and global consistency. In *Advances in neural information processing systems*, pages 321–328.
- Xiaojin Zhu, Zoubin Ghahramani, and John D Lafferty. 2003. Semi-supervised learning using Gaussian fields and harmonic functions. In *International conference on Machine learning*, pages 912–919.