

Continual Graph Convolutional Network for Text Classification

Tiandeng Wu^{1*}, Qijiong Liu^{2*}, Yi Cao¹, Yao Huang¹, Xiao-Ming Wu^{2†}, Jiandong Ding^{1†}

¹ Huawei Technologies Co., Ltd., China

² The Hong Kong Polytechnic University, Hong Kong
 {wutiandeng1, caoyi23, huangyao11, dingjiandong2}@huawei.com,
 jyonn.liu@connect.polyu.hk, xiao-ming.wu@polyu.edu.hk

Abstract

Graph convolutional network (GCN) has been successfully applied to capture global non-consecutive and long-distance semantic information for text classification. However, while GCN-based methods have shown promising results in offline evaluations, they commonly follow a *seen-token-seen-document* paradigm by constructing a fixed document-token graph and cannot make inferences on new documents. It is a challenge to deploy them in online systems to infer streaming text data. In this work, we present a continual GCN model (ContGCN) to generalize inferences from observed documents to unobserved documents. Concretely, we propose a new *all-token-any-document* paradigm to dynamically update the document-token graph in every batch during both the training and testing phases of an online system. Moreover, we design an occurrence memory module and a self-supervised contrastive learning objective to update ContGCN in a label-free manner. A 3-month A/B test on Huawei public opinion analysis system shows ContGCN achieves 8.86% performance gain compared with state-of-the-art methods. Offline experiments on five public datasets also show ContGCN can improve inference quality. The source code will be released at <https://github.com/Jyonn/ContGCN>.

Introduction

As one of the fundamental tasks in natural language processing, text classification has been extensively studied for decades and used in various applications (Xu et al. 2019; Abaho et al. 2021). In recent years, graph convolutional network (GCN) has been successfully applied in text classification (Yao, Mao, and Luo 2019; Lin et al. 2021) to capture global non-consecutive and long-distance semantic information such as token co-occurrence in a corpus.

A line of GCN-based methods (Li et al. 2019) perform document classification by simply constructing a homogeneous graph with each document as a node and modeling inter-document relations such as citation links, which however does not exploit document-token semantic information. Another line of GCN-based methods constructs heterogeneous document-token graphs, where each node represents

a document or a token, and each edge indicates a correlation factor between two nodes. However, they commonly follow a *seen-token-seen-document* (STSD) paradigm to construct a *fixed* document-token graph with all seen documents (labeled or unlabeled) and tokens and perform transductive inference. Given a new document with unobserved tokens, the trained model cannot make an inference because neither the document nor the unseen tokens are included in the graph. Hence, while these methods are effective in offline evaluations, they cannot be deployed in online systems to infer streaming text data.

To address this challenge, in this paper, we propose a new *all-token-any-document* (ATAD) paradigm to dynamically construct a document-token graph, and based on which we present a continual GCN model (**ContGCN**) for text classification. Specifically, we take the vocabulary of a pretrained language model (PLM) such as BERT (Devlin et al. 2019) as the global token set, so a new document can be tokenized into seen tokens from the vocabulary. We further form a document set which may contain any present documents (e.g., those in the current batch). The document-token graph then consists of tokens in the global token set and documents in the document set. The edge weights of the graph are dynamically calculated according to an occurrence memory module with historical token correlation information, and document embeddings are generated with pretrained semantic knowledge. In this way, ContGCN is enabled to perform inductive inference on streaming text data.

Furthermore, to address data distribution shift (Luo et al. 2022) which is prevalent in online services, we design a label-free online updating mechanism for ContGCN to save the cost and effort for periodical offline updates of the model with new text data. Specifically, we fine-tune the occurrence memory module according to the distribution shift of streaming text data and update the network parameters with a self-supervised contrastive learning objective.

ContGCN achieves favorable performance in both online and offline evaluations thanks to the proposed ATAD paradigm and label-free online updating mechanism. We have deployed ContGCN in an online text classification system – Huawei public opinion analysis system, which processes thousands of textual comments daily. A 3-month A/B test shows ContGCN achieves 8.86% performance gain compared with state-of-the-art methods. Offline evaluations

*Equal contribution (co-first authors). Author order determined by dice rolling.

†Corresponding authors.

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

on five real-world public datasets also demonstrate the effectiveness of ContGCN.

To summarize, our contributions are listed as follows:

- We propose a novel all-token-any-document paradigm and a continual GCN model to infer unobserved streaming text data, which, to our knowledge, is the first attempt to use GCN for online text classification.
- We design a label-free updating mechanism based on an occurrence memory module and a self-supervised contrastive learning objective, which enables to update ContGCN online with unlabeled documents.
- Extensive online A/B tests on an industrial text classification system and offline evaluations on five real-world datasets demonstrate the effectiveness of ContGCN.

Preliminaries

Graph Convolutional Network (GCN)

GCN (Welling and Kipf 2016) is a graph encoder that aggregates information from node neighborhoods. It is composed of a stack of graph convolutional layers. Formally, we use $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ to denote a graph, where $\mathcal{V} (n = |\mathcal{V}|)$ and \mathcal{E} are sets of nodes and edges, respectively. Note that each node $v \in \mathcal{V}$ is self-connected, i.e., $(v, v) \in \mathcal{E}$. We use $\mathbf{X} \in \mathbb{R}^{n \times d}$ to represent initial node representations, where d is the embedding dimension. To aggregate information from neighborhoods, a symmetric adjacency matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ is introduced, where A_{ij} is the correlation score of nodes v_i and v_j and $A_{ii} = 1$. The adjacency matrix is normalized as

$$\tilde{\mathbf{A}} = \mathbf{D}^{-\frac{1}{2}} \mathbf{A} \mathbf{D}^{\frac{1}{2}}, \quad (1)$$

where \mathbf{D} is a degree matrix and $D_{ii} = \sum_j A_{ij}$. At the k -th convolutional layers, the node embeddings are calculated as:

$$\mathbf{H}^{(k)} = \sigma \left(\tilde{\mathbf{A}} \mathbf{H}^{(k-1)} \mathbf{W}_k \right), \quad (2)$$

where $k \in \{1, 2, \dots, h\}$, h is the total number of convolutional layers, σ is the activation function, and $\mathbf{W}_k \in \mathbb{R}^{d \times d}$ is a trainable matrix. Specifically, $\mathbf{H}^{(0)} = \mathbf{X}$.

GCN-Based Text Classification

Text classification aims to classify documents into different categories. Formally, we use $\mathcal{D} (m = |\mathcal{D}|)$ to denote a set of given documents, which can be split into a training set $\mathcal{D}_{\text{train}}$ and a testing set $\mathcal{D}_{\text{test}}$. Each document \mathbf{s} can be represented as a list of tokens, $\mathbf{s} = (t_1^{(\mathbf{s})}, t_2^{(\mathbf{s})}, \dots, t_{|\mathbf{s}|}^{(\mathbf{s})})$, where $t_i^{(\mathbf{s})} \in \mathcal{T}$ is a token in the token vocabulary set $\mathcal{T} (u = |\mathcal{T}|)$.

Existing GCN-based text classification methods (Yao, Mao, and Luo 2019; Qiao et al. 2018; Lin et al. 2021) mainly follow a seen-token-seen-document paradigm to construct heterogeneous document-token graphs. Specifically, they first form a seen vocabulary set $\mathcal{T}_{\text{seen}}$ of size u' ($\mathcal{T}_{\text{seen}} \subset \mathcal{T}$), which contains all the seen tokens in the document set \mathcal{D} . Then, they construct a fixed document-token graph, whose nodes include all the seen tokens in $\mathcal{T}_{\text{seen}}$ and all the given documents in \mathcal{D} . The adjacency matrix \mathbf{A} of the graph is shown in Figure 1a, which consists of a token-token symmetric matrix $\mathbf{A}^{(1)} \in \mathbb{R}^{u' \times u'}$, a document-token matrix

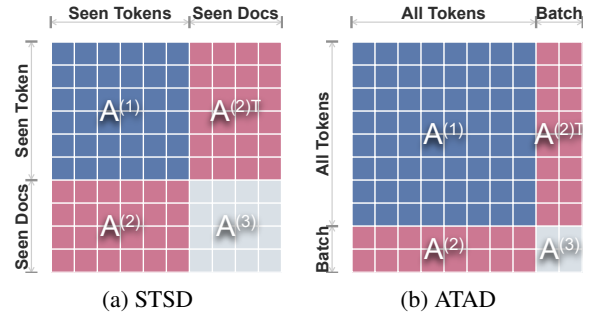


Figure 1: Comparison of the adjacency matrices. Left: seen-token-seen-document (STSD) paradigm (e.g., BertGCN). Right: proposed all-token-any-document (ATAD) paradigm.

$\mathbf{A}^{(2)} \in \mathbb{R}^{m \times u'}$, and a document-document identity matrix $\mathbf{A}^{(3)} \in \mathbb{R}^{m \times m}$. The adjacency matrix \mathbf{A} is fixed. Finally, GCN is applied to encode and classify the documents.

Method

The commonly adopted seen-token-seen-document (STSD) paradigm only allows to infer seen documents due to its transductive nature. To address this issue, we propose a novel all-token-any-document (ATAD) paradigm that leverages the global token vocabulary and dynamically updates the document-token graph to make inferences on unobserved documents.

Based on the ATAD paradigm, we propose a continual GCN model, namely ContGCN, as illustrated in Figure 2, which comprises of an adjacency matrix generator, a node encoder, and GCN encoders. Specifically, given a batch of input documents, the adjacency matrix generator updates the adjacency matrix based on the occurrence memory module and the current batch of documents, while the node encoder produces content-based node embeddings. The GCN encoder is then employed to capture the global-aware node representations. Finally, two training objectives are applied to train the ContGCN model.

Proposed All-Token-Any-Document Paradigm

In contrast to the seen-token-seen-document paradigm, which treats seen tokens ($\mathcal{T}_{\text{seen}}$) and seen documents (\mathcal{D}) as fixed graph nodes, our proposed all-token-any-document paradigm involves constructing a document-token graph with all tokens (\mathcal{T}) and dynamic documents (i.e., a batch of documents $B = \{\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_b\} \subset \mathcal{D}$ where b is the batch size).

Our ContGCN model is designed based on the all-token-any-document paradigm. We define all tokens as the vocabulary used for PLM tokenizers, allowing unseen words to be tokenized into sub-words that are already present in the vocabulary. When a new batch of data is fed into the model, the adjacency matrix and node embedding will be dynamically updated by the adjacency matrix generator and node encoder, respectively.

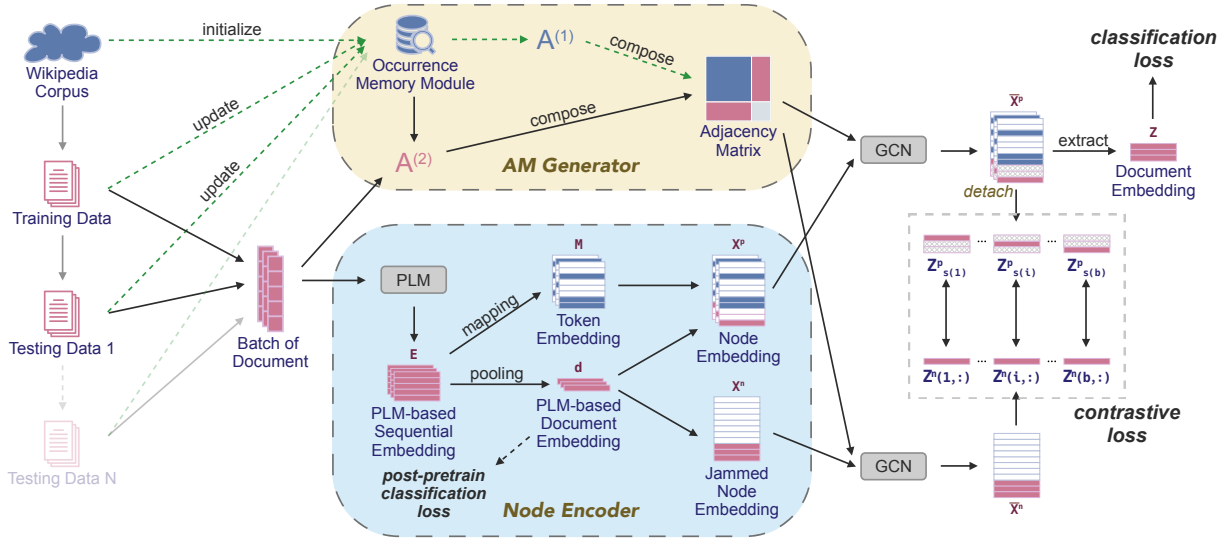


Figure 2: Framework of our ContGCN model. Green dotted lines represent operations before each phase of model training or testing. Two key components, i.e., AM Generator and Node Encoder, dynamically construct the adjacency matrix and generate node embeddings, which are then fed into a GCN encoder. Finally, our ContGCN model is trained with a classification loss and an anti-interference contrastive loss.

Adjacency Matrix Generator

As illustrated in Figure 1b, the adjacency matrix consists of three matrices: a token-token matrix $\mathbf{A}^{(1)} \in \mathbb{R}^{u \times u}$, a document-token matrix $\mathbf{A}^{(2)} \in \mathbb{R}^{b \times u}$, and a document-document identity matrix $\mathbf{A}^{(3)} \in \mathbb{R}^{b \times b}$.

The token-token matrix $\mathbf{A}^{(1)}$ is learned from the token occurrence knowledge of the corpus and is *phase-fixed*. This means that it will be refined when the model enters a new training or testing phase with the emergence of new corpora and token co-occurrence knowledge. The document-token matrix $\mathbf{A}^{(2)}$ is actively calculated based on the current batch of data and is used to update the adjacency matrix dynamically. Finally, the inner-document identity matrix $\mathbf{A}^{(3)}$ ensures that each document is not influenced by other samples during model learning or reasoning.

The Occurrence Memory Module (OMM) is a module that incrementally records historical statistics, which includes a document counter $s \in \mathbb{Z}^1$ that keeps track of the number of sentences, a token occurrence counter $\mathbf{c} \in \mathbb{Z}^u$ that records the number of sentences in which a token appears, and a token co-occurrence counter $\mathbf{C} \in \mathbb{Z}^{u \times u}$ that records the number of times two tokens appear in the same sentence. The OMM captures global non-consecutive semantic information and offers the following benefits: 1) it enables the dynamic calculation of the adjacency matrix for any batch of documents, and 2) it stores a large amount of previous general and domain-specific knowledge without the need for recalculation during updates. We develop a simple yet effective algorithm for updating the OMM, as described in Algorithm 1. As shown in Figure 2, the OMM is initialized with the Wikipedia corpus and updated by the training or testing data before model training or testing. Thus, $\mathbf{A}^{(1)}$ will be phasedly updated by PPMI (positive point-

wise mutual information), which is defined as:

$$\mathbf{A}_{i,j}^{(1)} = \begin{cases} 1, & \text{if } i = j, \\ \max\left(\log\left(s \frac{C_{i,j}}{c_{(i,:)}c_j}\right), 0\right), & \text{else.} \end{cases} \quad (3)$$

To obtain the document-token correlation for each document $s \in B$, we use TF-IDF (term frequency-inverse document frequency), which is calculated as:

$$\mathbf{A}_{s,t}^{(2)} = \frac{g(\mathbf{s}, t)}{|\mathbf{s}|} \log \frac{s}{c_t + 1}, \quad (4)$$

where $g(\mathbf{s}, t)$ represents the number of times token t appears in document s . The inner-document matrix $\mathbf{A}^{(3)}$ is an identity matrix, denoted as:

$$\mathbf{A}_{i,j}^{(3)} = \begin{cases} 1, & \text{if } i = j, \\ 0, & \text{else.} \end{cases} \quad (5)$$

Finally, the adjacency matrix \mathbf{A} can be composed by:

$$\mathbf{A} = \begin{pmatrix} \mathbf{A}^{(1)} & \mathbf{A}^{(2)\top} \\ \mathbf{A}^{(2)} & \mathbf{A}^{(3)} \end{pmatrix}. \quad (6)$$

Node Encoder

The effectiveness of pretrained language models (PLMs) such as BERT (Devlin et al. 2019), RoBERTa (Liu et al. 2019), and XLNet (Yang et al. 2019a) for text modeling has been widely demonstrated in various scenarios. Therefore, we utilize a PLM as a document encoder to capture semantic information for each document $s \in B$:

$$\mathbf{E}_{(s)} = \text{PLM}(s) \in \mathbb{R}^{l \times d}, \quad (7)$$

where l is the maximum document length for PLM, and each row of \mathbf{E} is a token embedding. We append special $\langle \text{PAD} \rangle$

Algorithm 1: Continual OMM updating algorithm

Input : A corpus or dataset \mathcal{D} , and OMM counters s , c , and C

Update document counter, i.e., $s \leftarrow s + \text{len}(\mathcal{D})$;
for each document in the corpus or dataset \mathcal{D} **do**
 for each sentence in the document **do**
 Update the count of token t_i if it appears in current sentence, i.e., $c[t_i] \leftarrow c[t_i] + 1$;
 for each token pair t_i and t_j in the sentence **do**
 if $t_i \neq t_j$ **then**
 Update the co-occurrence count, i.e.,
 $C[t_i][t_j] \leftarrow C[t_i][t_j] + 1$;
 end
 end
 end
end

Output: Updated OMM counters s , c , and C

tokens to short documents to align their length with other documents in the batch, following BERT. We average the hidden states of the first and last Transformer layers following previous works (Li et al. 2020; Su et al. 2021). We then perform an average pooling operation on $\mathbf{E}_{(s)}$ to obtain the unified document embedding $\mathbf{d}^{(s)} \in \mathbb{R}^d$. Following BertGCN (Lin et al. 2021), we form a batch-wise node embedding $\mathbf{X}^n \in \mathbb{R}^{(u+b) \times d}$:

$$\mathbf{X}^n = \left(\mathbf{0}, \dots, \mathbf{0}, \mathbf{d}^{(s_1)}, \dots, \mathbf{d}^{(s_b)} \right)^\top, \quad (8)$$

where $\mathbf{d}^{(s_j)}$ is the embedding of document j in the current batch B . However, the document embeddings will cause interference among one another when doing node message passing by GCN. To avoid interference, for each document j in the batch, we form a sample-specific node embedding $\mathbf{X}_{(s_j)}^p \in \mathbb{R}^{(u+b) \times d}$ by:

$$\mathbf{X}_{(s_j)}^p = \left(\mathbf{M}_{(s_j)}, \mathbf{0}, \dots, \mathbf{0}, \mathbf{d}^{(s_j)}, \mathbf{0}, \dots, \mathbf{0} \right)^\top, \quad (9)$$

where $\mathbf{M}_{(s_j)} \in \mathbb{R}^{u \times d}$ is a sample-specific token embedding matrix of document j :

$$\mathbf{M}_{(s_j)}(i, :) = \begin{cases} \mathbf{E}_{(s_j)}(k, :), & \text{if token } i \text{ of the vocabulary} \\ & \text{is the } k\text{-th token in } s_j, \\ \mathbf{0}, & \text{otherwise.} \end{cases} \quad (10)$$

We refer to \mathbf{X}^n as the *jammed* node embeddings for all documents in the current batch, and $\mathbf{X}_{(s)}^p$ as the *unjammed* node embedding of a single document s .

GCN Encoder

Once the adjacency matrix (\mathbf{A}) and node embeddings (unjammed $\mathbf{X}_{(s)}^p$ and jammed \mathbf{X}^n) are generated, the GCN encoder is applied to obtain graph-enhanced node representations. We denote the GCN-encoded unjammed and jammed

node representations as $\bar{\mathbf{X}}_{(s)}^p$ and $\bar{\mathbf{X}}^n$, respectively. Next, we extract the document embeddings from $\bar{\mathbf{X}}_{(s_j)}^p (\forall s_j \in B)$ to form $\mathbf{Z} \in \mathbb{R}^{b \times d}$ (see Figure 2), i.e.,

$$\mathbf{Z}(j, :) = \bar{\mathbf{X}}_{(s_j)}^p(j + u, :). \quad (11)$$

\mathbf{Z} will be used in the classification task. We then extract the document embeddings from $\bar{\mathbf{X}}_{(s_j)}^p$ to form $\mathbf{Z}_{(s_j)}^p \in \mathbb{R}^{b \times d}$ and those from $\bar{\mathbf{X}}^n$ to form $\mathbf{Z}^n \in \mathbb{R}^{b \times d}$ by:

$$\mathbf{Z}_{(s_j)}^p(i, :) = \bar{\mathbf{X}}_{(s_j)}^p(i + u, :) \text{ and} \quad (12)$$

$$\mathbf{Z}^n(i, :) = \bar{\mathbf{X}}^n(i + u, :), \quad (13)$$

as illustrated in Figure 2. $\mathbf{Z}_{(s_j)}^p$ and \mathbf{Z}^n will be used in the anti-interference contrastive task.

Training Objectives

To train the model, we employ two tasks: a document classification task and an anti-interference contrastive task.

The document classification task utilizes a multi-layer perceptron (MLP) classifier $f : \mathbb{R}^d \rightarrow \mathbb{R}^c$ (c is the number of document classes) with a softmax activation function to infer the probability distribution over all classes. The loss function can be defined as:

$$\mathcal{L}_{\text{cls}} = -\frac{1}{b} \sum_{j=1}^b \log \left(f(\mathbf{Z}(j, :))_{l_j} \right), \quad (14)$$

where l_j is the class label of the j -th document.

For the anti-interference contrastive task, the goal is to learn a representation space where semantically similar documents are closer to each other and dissimilar documents are farther apart. Hence, in the contrastive task, we enforce the GCN encoder to learn a mapping such that the jammed embedding of document j (i.e., $\mathbf{Z}^n(j, :)$) is close to its unjammed embedding (i.e., $\mathbf{Z}_{(s_j)}^p(j, :)$) while distant from the embeddings of other documents (i.e., $\mathbf{Z}_{(s_j)}^p(k, :)$, $k \neq j$) in the batch. The loss function can be calculated by:

$$\mathcal{L}_{\text{aic}} = -\frac{1}{b} \sum_{j=1}^b \log \left(\mathbf{y}_{(s_j)}(j) \right), \text{ where} \quad (15)$$

$$\mathbf{y}_{(s_j)} = \text{softmax} \left(\mathbf{Z}_{(s_j)}^p(\mathbf{Z}^n(j, :))^\top \right) \in \mathbb{R}^b. \quad (16)$$

The anti-interference contrastive task helps to learn representations robust to the interference between documents.

The overall loss function is a combination of the classification and contrastive tasks with a balancing parameter λ , denoted as:

$$\mathcal{L} = \mathcal{L}_{\text{cls}} + \lambda \mathcal{L}_{\text{aic}}. \quad (17)$$

Label-free Updating Mechanism (LUM). The occurrence memory module and the anti-interference contrastive task enables to continually update our ContGCN model with incoming unlabeled text data during inference. Hence, we name it label-free updating mechanism (LUM).

Model Training and Update

The training and updating of the ContGCN model involves three stages.

Stage 1: Before training. Prior to training, we perform post-pretraining on the pre-trained language model by using the classification task on the PLM-enhanced document embeddings. This pre-training helps to speed up the convergence of the model during the training process.

Stage 2: During training. During the training process, we use the multi-task training objective (Eq. 17) to train the ContGCN model.

Stage 3: During inference. When new test data is available, we first update the occurrence memory module using Algorithm 1. We then finetune the ContGCN model using the auxiliary anti-interference contrastive task (Eq. 15) to improve model performance during inference.

Experiments

Experimental Setups

Datasets. As described in (Lin et al. 2021), we have performed experiments on five text classification datasets which are commonly used in real-world applications: 20-News groups (20NG), Ohsumed, R52 Reuters, R8 Reuters, and Movie Review (MR) datasets. Table 1 presents the summarized statistics of these datasets. We randomly chose 10% of the training set for validation purposes for all datasets.

Baselines and Variants of Our Method. To validate the effectiveness of our proposed ContGCN model, we compare it with three types of state-of-the-art models: 1) traditional GCN-based models that do not utilize pretrained general semantic knowledge, including TextGCN (Yao, Mao, and Luo 2019) and TensorGCN (Liu et al. 2020); 2) transformer-based PLMs, including BERT (Devlin et al. 2019), RoBERTa (Liu et al. 2019) and XLNet (Yang et al. 2019b); 3) models that combine GCN with PLM, including TG-Transformer (Zhang and Zhang 2020), BertGCN (Lin et al. 2021) and RoBERTaGCN (Lin et al. 2021). As our ContGCN can be plugged by different PLMs, we adopt BERT, XLNet, and RoBERTa as alternatives, denoted as ContGCN_{BERT}, ContGCN_{XLNet}, and ContGCN_{RoBERTa}.

Implementation Details. We adopt the Adam optimizer (Kingma and Ba 2015) to train the network of our ContGCN model and the baseline models, which are consistent in the following parameters. The following parameters are kept consistent across all models: the number of graph convolutional layers, if applicable, is set to 3; the embedding dimension is set to 768; and the batch size is set to 64. In the post-pretraining phase, we set the learning rate for PLM to 1e-4. During training, we set different learning rates for PLM and other randomly initialized parameters including those of the GCN network, following Lin et al. (2021). Precisely, we set 1e-5 to finetune RoBERTa and BERT, 5e-6 to finetune XLNet, and 5e-4 to other parameters. We average the results of 10 runs as the final evaluation results.

Dataset	20NG	R8	R52	Ohsumed	MR
# Docs	18,846	7,674	9,100	7,400	10,662
# Training	11,314	5,485	6,532	3,357	7,108
# Test	7,532	2,189	2,568	4,043	3,554
# Classes	20	8	52	23	2
Avg. Length	221	66	70	136	20

Table 1: Dataset statistics.

Models	20NG	R8	R52	Ohsumed	MR
TextGCN	86.3	97.1	93.6	68.4	76.7
TensorGCN	87.7	98.0	95.0	70.1	77.9
BERT	85.3	97.8	96.4	70.5	85.7
RoBERTa	83.8	97.8	96.2	70.7	89.4
XLNet	85.1	98.0	<u>96.6</u>	70.7	87.2
TG-Transformer	-	98.1	95.2	70.4	-
BertGCN	89.3	98.1	<u>96.6</u>	<u>72.8</u>	86.0
RoBERTaGCN	<u>89.5</u>	<u>98.2</u>	96.1	<u>72.8</u>	<u>89.7</u>
ContGCN _{BERT}	89.4	98.3	96.9	73.1	86.4
ContGCN _{XLNet}	89.7	98.5	97.0	73.1	88.7
ContGCN _{RoBERTa}	90.1	98.6	96.6	73.4	91.3

Table 2: Comparison of ContGCN with state-of-the-art models in offline evaluation. The best results are in bold-face, and the second best results are underlined.

Offline Evaluation

We conduct an offline evaluation of our ContGCN model and state-of-the-art baselines on five datasets. Table 2 presents the overall performance of all methods, and the following observations can be made. **First**, PLM-only methods mostly outperform GCN-only methods due to their pre-learned semantic knowledge. However, GCN-only methods can build more document-token edges for better semantic comprehension on datasets with long documents, such as 20NG, but they struggle on datasets with short documents, such as MR. **Second**, PLM-empowered GCN methods combine the strengths of both PLMs and GCNs and outperform both PLM-only and GCN-only methods. **Third**, our ContGCN model achieves state-of-the-art performance on five datasets, thanks to 1) the proposed all-token-any-document paradigm that leverages general semantic knowledge from a large Wikipedia corpus and 2) the proposed contrastive learning objective that reduces inter-document interference. Notably, our ContGCN_{RoBERTa} achieves the best performance on four datasets.

Ablation Study

First, we study the effect of different components of ContGCN, including Wikipedia initialization, OMM updating, and anti-interference contrastive task on offline performance. Based on the results from Table 3, we can conclude the following. **First**, on the 20NG dataset, we observed that Wikipedia initialization is less effective, which is likely because the lengthy documents already contain sufficient non-consecutive knowledge during OMM updating. Apart from this, removing any of these components

Models	20NG	R8	Ohsumed
ContGCN _{RoBERTa}	90.1	98.6	73.4
w/o Wikipedia Init	89.9	98.2	73.1
w/o OMM Updating	89.6	98.3	73.0
w/o Contrastive Loss	89.7	98.5	73.2
ContGCN _{XLNet}	89.7	98.5	73.1
w/o Wikipedia Init	89.8	98.3	72.8
w/o OMM Updating	89.4	98.2	72.7
w/o Contrastive Loss	89.5	98.2	73.0

Table 3: Influence of Wikipedia initialization, OMM updating, and the anti-interference contrastive task.

Variants	1/6	2/6	3/6	4/6	5/6	6/6
ContGCN*	86.4	87.3	88.1	88.6	89.0	89.6
ContGCN	86.3	87.1	87.8	88.2	88.7	89.1
ContGCN ^α	86.1	86.9	87.5	87.9	88.3	88.7
ContGCN ^β	86.0	86.2	86.4	86.6	86.9	87.1

Table 4: Comparisons of variants of ContGCN_{RoBERTa} in the online learning scenario on the 20NG dataset. ContGCN* is retrained from scratch in each session with all previously seen data. ContGCN^α is updated without the contrastive loss. ContGCN^β is updated without LUM.

leads to a decline in performance for both ContGCN_{RoBERTa} and ContGCN_{XLNet} on three datasets, confirming their effectiveness. **Second**, models without OMM updating show the worst performance, indicating the importance of non-consecutive semantic information in training and testing.

Next, we study the updating strategy of ContGCN in the online learning scenario. As shown in Table 4, we can make the following observations. **First**, by comparing ContGCN with ContGCN^α and ContGCN^β, it can be seen that both OMM updating and the contrastive loss are effective. **Second**, compared with retraining from scratch, i.e., ContGCN*, ContGCN requires less computational resources and time to update yet still achieves competitive performance.

Impact of Anti-interference Contrastive Learning

Here, we study the balancing parameter λ which weights the auxiliary anti-interference contrastive loss. We conduct experiments on the 20NG, R8 and Ohsumed datasets with ContGCN_{RoBERTa} model, varying λ in $\{0.001, 0.01, 0.02, 0.03, 0.05, 0.10, 0.20\}$. As demonstrated in Figure 3, we can make the following observations. **First**, the reliance on the auxiliary task varies for different datasets. Specifically, the model achieves the best performance on the 20NG, R8, and Ohsumed datasets when λ is set to 0.03, 0.02, and 0.05, respectively. **Second**, for each dataset, as λ increases, the performance first increases and then decreases. Hence, it is essential to select a good λ .

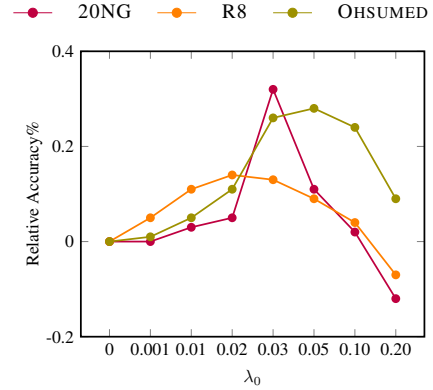


Figure 3: Influence of the parameter λ that weights the anti-interference contrastive loss. *Relative accuracy (%)* means the difference between the accuracy achieved with $\lambda = \lambda_0$ and that achieved with $\lambda = 0$.

Online Evaluation

Fixed Testing Data. Figure 4 illustrates the performance of our model and baselines in an online learning scenario where the training/updating data is incremental and the testing data is fixed. Based on the results, we can draw the following observations: **First**, GCN-based methods such as TextGCN and RoBERTaGCN cannot be updated and are incapable of reasoning about unobserved data as they construct fixed graphs based on the original corpus. Therefore, their performance remains constant over time, as shown by the dotted lines in Figure 4a. **Second**, as the updating data increases, the performance of all updatable models improves. **Third**, our proposed ContGCN method outperforms all baselines in all sessions. **Fourth**, due to the limitations of STSD-based GCN methods, we introduce a RoBERTaGCN_{scratch} model that retrains from scratch in each session with all previously seen data. However, this model still falls short compared to ContGCN_{RoBERTa} due to the utilization of Wikipedia Initialization and the unjammed node embedding in ContGCN_{RoBERTa}. Moreover, as illustrated in Figure 4b, the updating time of RoBERTaGCN_{scratch} increases almost linearly w.r.t. data size, making it unsuitable for online learning. The finetuning time ratios of ContGCN and RoBERTa are closer to or less than 1, indicating that for these models, each session takes up a comparable amount of time.

Fixed Training Data. We have deployed ContGCN on Huawei public opinion analysis system, under the scenario that the training data is fixed and the updating/testing data is incremental. We use an optimized variant of RoBERTa (Liu et al. 2019) – RoBERTa_{wm-ext} (Xu 2021), tailored for Chinese text classification, which will still be referred to as RoBERTa below and in Table 5. We implement RoBERTaGCN by plugging RoBERTa into BertGCN (Lin et al. 2021), which cannot infer unobserved documents due to the STSD scheme and hence the results are unavailable in the following months. We then implement ContGCN based on RoBERTa and the proposed ATAD scheme, i.e.,

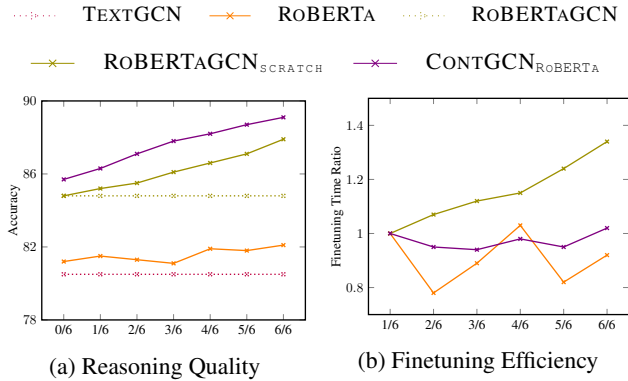


Figure 4: Comparison between our ContGCN model and baselines in an online learning scenario. We divide the 20NG dataset into training, testing, and updating sets by the ratio of 2:2:6. We trained each model with the training set to learn an initial version. Then, we divided the updating set into six equal parts and gradually fed each part to the model for finetuning. The *finetuning time ratio* in (b) is calculated by the finetuning time of the current session over that of the first session. For each training or updating session, we used 10% of the training set as the validation set.

Models	0th	1st	2nd	3rd
RoBERTaGCN	91.7	N/A	N/A	N/A
RoBERTa	87.6	86.8	85.2	83.5
ContGCN _{RoBERTa} ^β	92.8	90.3	89.9	88.2
ContGCN _{RoBERTa}	92.8	92.5	92.0	90.9

Table 5: Comparison of our ContGCN model with RoBERTa in an industrial online learning scenario. All models are first trained offline (in the 0th month) with a labeled dataset. After deployed, ContGCN_{RoBERTa} performs online learning with LUM. ContGCN_{RoBERTa}^β is a static network with parameters fixed after the initial training.

ContGCN_{RoBERTa}. After RoBERTa and ContGCN_{RoBERTa} were trained offline initially (i.e., in the 0th month), we deploy them online for comparison. As illustrated in Table 5, our ContGCN_{RoBERTa} has achieved 5.94%, 6.57%, 7.98%, and 8.86% gains in accuracy over RoBERTa in the 0th, 1st, 2nd, and 3rd month respectively. Besides, due to the *distribution shift* of public opinions, the accuracy drops slightly over time. However, ContGCN_{RoBERTa} still maintains higher performance than RoBERTa. Furthermore, by removing the label-free update mechanism (i.e., ContGCN_{RoBERTa}^β), the performance drops significantly, which demonstrates the capability of our ContGCN model in continual learning.

Related Work

Graph Convolutional Network. Graph Convolutional Networks (Welling and Kipf 2016) (GCNs) have become increasingly popular in recent years due to their ability to capture the structural relations among data (Hamilton, Ying,

and Leskovec 2017; Li, Han, and Wu 2018). They can learn representations of graph data by propagating information between nodes in the graph. The popularity of GCNs can be attributed to their versatility and effectiveness in various applications (Zhang et al. 2019), including image classification (Hong et al. 2020), video understanding (Huang et al. 2020), social recommendation (Fan et al. 2019), and text classification (Yao, Mao, and Luo 2019; Lin et al. 2021).

Text Classification. Text classification is a critical and fundamental task in the field of natural language processing. Early studies (Jacovi, Sar Shalom, and Goldberg 2018; Sari, Rini, and Malik 2019) utilize word embedding methods (Mikolov et al. 2013; Pennington, Socher, and Manning 2014) or apply traditional models (Zhang, Zhao, and LeCun 2015; Lai et al. 2015) such as convolutional neural networks (Krizhevsky, Sutskever, and Hinton 2012) or recurrent neural networks (Rumelhart, Hinton, and Williams 1985) to learn textual knowledge. In recent years, Transformer-based pretrained language models such as BERT (Devlin et al. 2019) have been introduced (Sun et al. 2019) for text classification due to its strong ability of semantic comprehension. However, these models do not effectively utilize global semantic information such as token co-occurrence in a corpus.

GCN-based Text Classification. GCNs have been gaining attention in text classification, owing to their ability to model non-structured data and capture global dependencies, such as high-order neighborhood information (Yao, Mao, and Luo 2019; Lin et al. 2021). Unlike some GCN-based methods (Li et al. 2019; Xie et al. 2021; Li et al. 2021) that construct homogeneous document graphs, we follow Yao, Mao, and Luo (2019) to construct a document-token graph to better capture semantic relations. In contrast to previous methods that typically construct fixed graphs and are limited in their ability to reason about unobserved documents, our proposed ContGCN model employs an all-token-any-document paradigm, enabling making inferences on new data in online systems.

Conclusion

To deploy GCN-based text classification methods to online industrial systems, we propose a ContGCN model with a novel all-token-any-document paradigm and a label-free updating mechanism, which endow the model the ability to infer unobserved documents and enable to continually update the model during inference time. To our knowledge, this is the first attempt to use GCN for online text classification. Extensive online and offline evaluations validate the effectiveness of ContGCN, which achieves favorable performance compared with various state-of-the-art methods.

Acknowledgments

The authors would like to thank Qimai Li for valuable discussion and the anonymous reviewers for their helpful comments. Q. Liu and X. Wu were supported by GRF No.15222220 funded by the UGC and ITS/359/21FP funded by the ITC of Hong Kong.

References

- Abaho, M.; Bollegala, D.; Williamson, P.; and Dodd, S. 2021. Detect and Classify – Joint Span Detection and Classification for Health Outcomes. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 8709–8721. Online and Punta Cana, Dominican Republic: Association for Computational Linguistics.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 4171–4186. Minneapolis, Minnesota: Association for Computational Linguistics.
- Fan, W.; Ma, Y.; Li, Q.; He, Y.; Zhao, E.; Tang, J.; and Yin, D. 2019. Graph Neural Networks for Social Recommendation. In *The world wide web conference*, 417–426.
- Hamilton, W.; Ying, Z.; and Leskovec, J. 2017. Inductive Representation Learning on Large Graphs. *Advances in neural information processing systems*, 30.
- Hong, D.; Gao, L.; Yao, J.; Zhang, B.; Plaza, A.; and Chanussot, J. 2020. Graph Convolutional Networks For Hyperspectral Image Classification. *IEEE Transactions on Geoscience and Remote Sensing*, 59(7): 5966–5978.
- Huang, D.; Chen, P.; Zeng, R.; Du, Q.; Tan, M.; and Gan, C. 2020. Location-aware Graph Convolutional Networks for Video Question Answering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 11021–11028.
- Jacovi, A.; Sar Shalom, O.; and Goldberg, Y. 2018. Understanding Convolutional Neural Networks for Text Classification. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, 56–65. Brussels, Belgium: Association for Computational Linguistics.
- Kingma, D. P.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. *International Conference on Learning Representations*.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. ImageNet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25.
- Lai, S.; Xu, L.; Liu, K.; and Zhao, J. 2015. Recurrent Convolutional Neural Networks for Text Classification. In *Proceedings of the AAAI conference on artificial intelligence*, volume 29.
- Li, B.; Zhou, H.; He, J.; Wang, M.; Yang, Y.; and Li, L. 2020. On the Sentence Embeddings from Pre-trained Language Models. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 9119–9130. Online: Association for Computational Linguistics.
- Li, Q.; Han, Z.; and Wu, X.-M. 2018. Deeper Insights into Graph Convolutional Networks for Semi-supervised Learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32.
- Li, Q.; Wu, X.-M.; Liu, H.; Zhang, X.; and Guan, Z. 2019. Label Efficient Semi-supervised Learning via Graph Filtering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9582–9591.
- Li, Q.; Zhang, X.; Liu, H.; Dai, Q.; and Wu, X.-M. 2021. Dimensionwise Separable 2-D Graph Convolution for Unsupervised and Semi-supervised Learning on Graphs. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 953–963.
- Lin, Y.; Meng, Y.; Sun, X.; Han, Q.; Kuang, K.; Li, J.; and Wu, F. 2021. BertGCN: Transductive Text Classification by Combining GNN and BERT. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, 1456–1462.
- Liu, X.; You, X.; Zhang, X.; Wu, J.; and Lv, P. 2020. Tensor Graph Convolutional Networks for Text Classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 8409–8416.
- Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; and Stoyanov, V. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. *CoRR*, abs/1907.11692.
- Luo, R.; Sinha, R.; Hindy, A.; Zhao, S.; Savarese, S.; Schmerling, E.; and Pavone, M. 2022. Online Distribution Shift Detection via Recency Prediction. *arXiv preprint arXiv:2211.09916*.
- Mikolov, T.; Sutskever, I.; Chen, K.; Corrado, G. S.; and Dean, J. 2013. Distributed Representations of Words and Phrases and Their Compositionality. In *Advances in neural information processing systems*, 3111–3119.
- Pennington, J.; Socher, R.; and Manning, C. D. 2014. Glove: Global Vectors for Word Representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 1532–1543.
- Qiao, C.; Huang, B.; Niu, G.; Li, D.; Dong, D.; He, W.; Yu, D.; and Wu, H. 2018. A New Method of Region Embedding for Text Classification. In *International Conference on Learning Representations*.
- Rumelhart, D. E.; Hinton, G. E.; and Williams, R. J. 1985. Learning Internal Representations by Error Propagation. Technical report, California Univ San Diego La Jolla Inst for Cognitive Science.
- Sari, W.; Rini, D.; and Malik, R. 2019. Text Classification Using Long Short-Term Memory With GloVe Features. *Jurnal Ilmiah Teknik Elektro Komputer dan Informatika*, 5(2): 85–100.
- Su, J.; Cao, J.; Liu, W.; and Ou, Y. 2021. Whitening Sentence Representations for Better Semantics and Faster retrieval. *arXiv preprint arXiv:2103.15316*.
- Sun, C.; Qiu, X.; Xu, Y.; and Huang, X. 2019. How to Fine-tune BERT for Text Classification? In *Chinese Computational Linguistics: 18th China National Conference, CCL 2019, Kunming, China, October 18–20, 2019, Proceedings 18*, 194–206. Springer.

Welling, M.; and Kipf, T. N. 2016. Semi-supervised Classification with Graph Convolutional Networks. In *J. International Conference on Learning Representations (ICLR 2017)*.

Xie, Q.; Huang, J.; Du, P.; Peng, M.; and Nie, J.-Y. 2021. Inductive Topic Variational Graph Auto-Encoder for Text Classification. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 4218–4227. Online: Association for Computational Linguistics.

Xu, G.; Meng, Y.; Qiu, X.; Yu, Z.; and Wu, X. 2019. Sentiment Analysis of Comment Texts Based on BiLSTM. *Ieee Access*, 7: 51522–51532.

Xu, Z. 2021. RoBERTa-wwm-ext Fine-Tuning for Chinese Text Classification. *arXiv preprint arXiv:2103.00492*.

Yang, Z.; Dai, Z.; Yang, Y.; Carbonell, J.; Salakhutdinov, R. R.; and Le, Q. V. 2019a. XLNet: Generalized Autoregressive Pretraining for Language Understanding. In *Advances in Neural Information Processing Systems*, volume 32.

Yang, Z.; Dai, Z.; Yang, Y.; Carbonell, J.; Salakhutdinov, R. R.; and Le, Q. V. 2019b. Xlnet: Generalized Autoregressive Pretraining for Language Understanding. *Advances in neural information processing systems*, 32.

Yao, L.; Mao, C.; and Luo, Y. 2019. Graph Convolutional Networks for Text Classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 7370–7377.

Zhang, H.; and Zhang, J. 2020. Text Graph Transformer for Document Classification. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*.

Zhang, S.; Tong, H.; Xu, J.; and Maciejewski, R. 2019. Graph Convolutional Networks: A Comprehensive Review. *Computational Social Networks*, 6(1): 1–23.

Zhang, X.; Zhao, J.; and LeCun, Y. 2015. Character-level Convolutional Networks for Text Classification. In *Advances in neural information processing systems*, 649–657.