

Weakly Supervised Text Classification using Supervision Signals from a Language Model

Ziqian Zeng^{1,2}, Weimin Ni¹, Tianqing Fang², Xiang Li³, Xinran Zhao⁴, and Yangqiu Song²

¹Shien-Ming Wu School of Intelligent Engineering, South China University of Technology, China

²Department of CSE, HKUST, Hong Kong, China

³Department of Computer Science, University of Illinois, Urbana Champaign, US

⁴Computer Science Department, Stanford University, US

zqzeng@scut.edu.cn, mewmn@mail.scut.edu.cn, xiangl12@illinois.edu
xzhaoar@stanford.edu, {tfangaa, yqsong}@cse.ust.hk

Abstract

Solving text classification in a weakly supervised manner is important for real-world applications where human annotations are scarce. In this paper, we propose to query a masked language model with cloze style prompts to obtain supervision signals. We design a prompt which combines the document itself and “this article is talking about [MASK].” A masked language model can generate words for the [MASK] token. The generated words which summarize the content of a document can be utilized as supervision signals. We propose a latent variable model to learn a word distribution learner which associates generated words to pre-defined categories and a document classifier simultaneously without using any annotated data. Evaluation on three datasets, AG-News, 20NewsGroups, and UCI-News, shows that our method can outperform baselines by 2%, 4%, and 3%.

1 Introduction

Text classification is a fundamental task in Natural Language Processing (NLP) with diverse real-world applications such as identifying relevant documents of a case in legal proceedings (Roitblat et al., 2010), and classifying victim’s requests (e.g., food, shelter, and medical aids) on social media platforms during earthquakes (Caragea et al., 2011). Current state-of-the-art text classification methods (Zhang et al., 2015; Zhou et al., 2016; Johnson and Zhang, 2017) still need a large number of annotated data. However, in the real world, naturally annotated data are rare and human annotations are expensive. Solving the text classification task without using annotated data but exploiting inexpensive supervision signals is worth investigation.

In the weakly supervised setting, any annotated document is not accessible, but inexpensive supervision signals such as label surface names or keywords can be used. Existing weakly supervised text classification methods (Meng et al., 2018, 2019;

Mekala and Shang, 2020; Meng et al., 2020) first used seed keywords to retrieve more keywords, and then created pseudo labels for documents and then train a model in a “standard” supervised learning manner. In previous work, supervision signals are restricted to a small set of keywords from documents contents.

Recent work shows that prompts can probe knowledge from PLMs (Devlin et al., 2019; Radford et al., 2019) and the knowledge can provide supervision signals to solve different NLP tasks including relation extraction (Shin et al., 2020), question answering (Petroni et al., 2020), and summarization (Radford et al., 2019). For example, (Petroni et al., 2019) solved the knowledge base completion task by querying an MLM with a prompt “Alan Turing was born in [MASK].” Using prompts to generate supervision signals for text classification is worth exploring.

We propose to query an MLM with a prompt which combines the document itself and “this article is talking about [MASK].”, and use generated words for the [MASK] token as supervision signals. For example, in Figure 1, given a prompt “The radio telescope at arecibo observatory will begin mapping the known galaxy on friday, scientists said. This article is talking about [MASK].”, an MLM predicts “astronomy”, “galaxies”, “radio”, “science”, and “galaxy” for the [MASK] token. These words summarize the topic of the document. Hence, they can be used as supervision signals. Besides generating signal words, an intuitive approach to obtain supervision signals is by extracting important words from documents. We will compare two types of supervision signals.

After obtaining signal words, we need to associate these words to pre-defined categories. We propose a latent variable model (WDDC) to learn a Word Distribution learner and a Document Classifier simultaneously without using any annotated data. A word distribution learner aims to learn

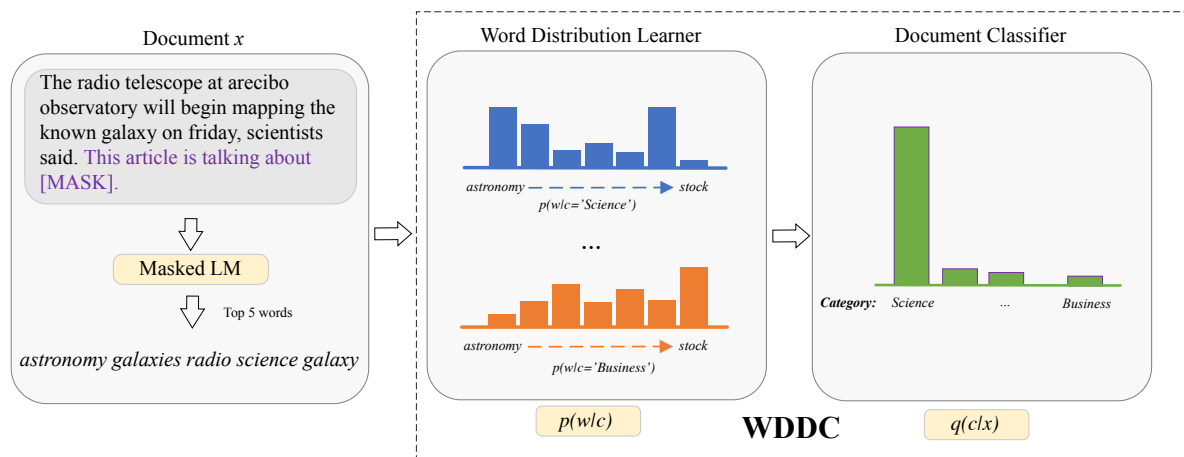


Figure 1: We combine a document and a cloze style sentence “This article is talking about [MASK]” to query a masked LM. It generates a set of words for the [MASK] token. These words are likely to summarize the topic of a document. After obtaining words such as “astronomy” and “galaxies”, human beings can easily infer that this article is talking *science* rather than *business* because we know these words are frequently used in *science* topic. Word distributions given pre-defined categories bridge supervision signals (generated words) and our goal (the category of a document). The proposed model (WDDC) can learn word distributions given pre-defined categories and a document classifier simultaneously.

a probability of a generated word w given a category c , i.e., $p(w|c)$. A document classifier aims to learn a probability of a category c given a document x , i.e., $p(c|x)$. These two goals could be optimized simultaneously via maximizing the log-likelihood of generated words by introducing the category as a latent variable. In our latent variable model, a word distribution learner and a document classifier can be parameterized by any neural network.

Our contributions are summarized as follows,

- We propose to query an MLM using a prompt which combines the document and a cloze style sentence “this article is talking about [MASK]”. We use generated words for the [MASK] token as supervision signals in the weakly supervised text classification task.

- We propose a latent variable model (WDDC) to learn a word distribution learner which associates generated words to pre-defined categories and a document classifier simultaneously without using any annotated data.

- The experimental results show that the proposed method WDDC can outperform other weakly supervised baselines in three datasets.

The code is available at <https://github.com/HKUST-KnowComp/WDDC>.

2 Related Work

In this section, we review the related work on querying an MLM with prompts, weakly supervised text

classification, zero-shot text classification, and variational methods.

Querying an MLM with Prompts. Querying an MLM with cloze style prompts provides a new direction to solve some NLP tasks in an unsupervised manner. (Petroni et al., 2019) queried an MLM using manually designed prompts to solve a knowledge base completion task. For example, in order to complete the missing entity x in (Alan Turing, born in, x), they designed a prompt “Alan Turing was born in [MASK]” to query an MLM. The generated word for the [MASK] token can be directly used to complete the missing fact. By querying language models, some NLP tasks such as relation extraction (Shin et al., 2020), question answering (Radford et al., 2019; Petroni et al., 2020), summarization (Radford et al., 2019) could be solved in an unsupervised manner. However, not all NLP tasks can directly use generated words from an MLM in downstream tasks. Some tasks such as sentiment analysis and textual entailment (Shin et al., 2020) need more steps for inference. For example, in the sentiment analysis task, (Shin et al., 2020) used annotated data to train a classifier that links generated words to pre-defined categories. Our work does not require annotated data for inference.

Weakly Supervised Text Classification. In the weakly supervised text classification task, any labeled documents are not allowed, but label sur-

face names or limited word-level descriptions of each category can be used. Dataless (Chang et al., 2008; Song and Roth, 2014) used Explicit Semantic Analysis (ESA) vectors (Gabrilovich et al., 2007) to represent label name and documents. Predictions are based on the label-document similarity. Recently, (Meng et al., 2018, 2019; Mekala and Shang, 2020; Meng et al., 2020; Schick and Schütze, 2021; Schick and Schütze, 2021; Zhang et al., 2022) trained neural text classifiers in a weakly supervised manner. They generated pseudo labels for documents to pre-train a neural classifier and then performed self-training on unlabeled data for model refinement. LOTClass (Meng et al., 2020) is relevant to our work because they also used pre-trained language models. They used a LM to retrieve a set of semantically correlated words for each class, and then fine-tuned the LM to predict these words. Finally, they performed self-training on unlabeled data. Our work is different from LOTClass because we obtain supervision signals by querying an MLM with cloze style prompts and we propose a latent variable model to learn document classifier rather than using the self-training procedure. PRBOOST (Zhang et al., 2022) is also relevant to our work because they also use prompts to generate weak labels. PRBOOST first generated rules by using a small amount of labeled data, then asked human annotators to select high-quality rules to generate weak labels. Finally, they trained a new model in a self-training manner. Our work is different from PRBOOST because our method associates predicted words with labels in an unsupervised manner while PRBOOST maps prompting based rules to labels by involving human feedback.

Zero-Shot Text Classification. In zero-shot learning settings, the classes covered by training instances and the classes we aim to classify are disjoint. Zero-shot learning text classification methods (Xia et al., 2018; Rios and Kavuluru, 2018; Zhang et al., 2019; Liu et al., 2019) generalized seen classes to unseen classes by learning semantic relationships between classes and documents via embeddings or semantic knowledge sources. However, zero-shot learning still requires annotated data for the seen classes training. We cannot apply zero-shot learning methods to weakly supervised settings where no annotated document is available.

Variational Methods. Variational autoencoders (Kingma and Welling, 2014; Rezende et al., 2014)

consists of an encoder and a decoder. The encoder estimates posterior probabilities and the decoder estimates the reconstruction likelihood given a latent variable. The objective function is to maximize the reconstruction likelihood of the observed variable. The latent variable in VAEs is continuous variable. Recently, many research works (Titov and Khodam, 2015; Marcheggiani and Titov, 2016; Šuster et al., 2016; Zhang et al., 2018; Chen et al., 2018; Zeng et al., 2019; Liang et al., 2019) use VAEs to solve different NLP tasks such as relation discovery, question answering, sentiment classification, etc. In above works, the latent variables are discrete variables. For example, (Marcheggiani and Titov, 2016) aimed to solve unsupervised open-domain relation discovery. The objective function is to reconstruct the likelihood of two entities. They introduced relation as the latent variable. The encoder is a relation classifier, which predicts a semantic relation between two entities. The decoder reconstructs entities given the predicted relation. Our method is also based on VAEs with a discrete latent variable but the estimated probabilities and the objective function are different.

3 Methodology

In this section, we first introduce how to obtain supervision signals from an MLM and document itself, and then we introduce a latent variable model to learn a word distribution learner and a document classifier simultaneously.

3.1 Supervision Signals

3.1.1 Signal Words

Given a document, our goal is to obtain topic relevant words which are used as supervision signals. To achieve this, we append a cloze style sentence to the document at the end as a prompt. A prompt is designed as “[CLS] + document + This article is talking about [MASK]. + [SEP].” The [MASK] token serves as a placeholder for a topic relevant word which can summarize the document. It mimics the reading comprehension task which is using a word to summarize the content of a document. We select top k generated words as supervision signals.

Instead of generating signal words, a natural way to obtain supervision signals is by extracting words from the document. To achieve this, we extract all nouns and proper nouns in the document using part-of-speech tagger (Kristina et al., 2003). Since

Table 1: Signal words from an MLM and from the document(Doc).

Text	Label	Signal Words
The world ’s top two players roger federer and andy roddick reached the semifinals friday at the thailand open.	Sports	MLM: tennis, thailand, federer, seeds, wimbledon Doc: world, players, federer, andy, roddick, semifinals, friday, thailand, open
These circuits abound in most electronic project books. It has LED indicators also.	Science	MLM: circuits, computers, electronics, computing, graphs Doc: circuits, project, books, LED, indicators
Scientists discover a genetic indicator that could help prevent suicides.	Health	MLM: suicide, genetics, cancer, hiv, health Doc: scientists, indicator, suicides

most of the generated words from an MLM are nouns and proper nouns, so we only extract words with two types of part-of-speech.

Table 1 shows top 5 predictions from an MLM (BERT (Devlin et al., 2019)) given prompts and extracted nouns and proper nouns from documents. For the first document, an MLM can infer that it is talking about a tennis match although “tennis” does not appear in the document. It also generates some relevant words such as “wimbledon.” In this case, the MLM is better than extraction. For the second document, the first word from the MLM precisely summarizes the document. However, the MLM also generates a few words which are related to *computer*. Unfortunately *computer* is also a category in this dataset. Compared to the MLM, the extracting way is safer in this case. For the third document, an MLM generates “health” which is an exact match of the label surface name although “cancer” and “hiv” are not faithful to the original document. We will evaluate generation and extraction methods in the experiment.

3.1.2 Remove Non-discriminated Words

Words generated from an MLM are not always category discriminated. Non-discriminated words can harm the performance of inference. The intuition of removing non-discriminated words is that if some words appear in different categories with similar frequency, then it is possible that these words are not category-discriminated. Since we cannot access labels, the label in the following computation means the pseudo label. The pseudo label generation process is shown in section 3.1.3. Inspired by category-indicative measurement, (Mekala and Shang, 2020), we define category-indicative index:

$$CII(c_i, w) = \frac{f(c_i, w)}{f(c_i)}, \quad (1)$$

where $f(c_i, w)$ is the number of occurrences of the signal word w in the documents which are labeled as c_i , and $f(c_i)$ is the total number of occurrences of all signal words in documents which are labeled as c_i .

We define category-indicative ratio as,

$$CIR(w) = \frac{CII(c_i, w)}{CII(c_j, w)}, \quad (2)$$

where $CII(c_i, w)$ is the maximum value among all categories, $CII(c_j, w)$ is the second maximum value all categories. Larger value of $CIR(w)$ indicates w is more discriminated. If $CII(c_j, w)$ is equal to 0, we will assign a large value to $CIR(w)$. If $CIR(w) < t$, we consider w is not discriminated and we remove w from signal words set.

3.1.3 Pseudo Label Generation

We assign pseudo labels to data based on label-word similarity. We represent a word using static representation which is introduced by (Mekala and Shang, 2020). Given a word w , static representation $\mathcal{SR}(w)$ is computed by averaging the contextualized embeddings of all its occurrences in the corpus. The label-word similarity is the cosine similarity between the static representation of the label surface name and the static representation of signal words. If the label surface name or the supervision signal contains more than one word, we take the average of the static representations of all words. We assign a sample with the pseudo label which yields the maximum similarity value among all classes. And the similarity value should be greater

than a threshold γ . Setting a threshold can result in more accurate pseudo label assignments although the size of pseudo labeled data will shrink.

To summarize, there are three steps to obtain clean signal words: (1) Obtain signal words from an MLM or a document. (2) Generate pseudo labels. (3) Remove signal words which have low category-indicative ratio values.

3.2 Model Training

After getting clean signal words, we then propose a latent variable model to learn a word distribution learner and a document classifier simultaneously.

Since there is no annotated data available, in order to best explain the observed data, i.e., signal words, the objective of our model is to maximize the log-likelihood of signal words. The ultimate goal is to identify the category of a document, hence, we introduce a latent variable C representing the category, into the objective function. Further, by applying Jensen’s inequality (Jensen et al., 1906), we can derive an evidence lower bound (ELBO) of the log-likelihood. We define the objective function as follows,

$$\begin{aligned}
\mathcal{L}_o &= \sum_{x \in X} \sum_{w_r \in \mathcal{R}_x} \log p(w_r) \\
&= \sum_{x \in X} \sum_{w_r \in \mathcal{R}_x} \log \sum_c p(w_r, c) \\
&= \sum_{x \in X} \sum_{w_r \in \mathcal{R}_x} \log \sum_c q(c|x) \left[\frac{p(w_r, c)}{q(c|x)} \right] \\
&\geq \sum_{x \in X} \sum_{w_r \in \mathcal{R}_x} \sum_c q(c|x) \left[\log \frac{p(w_r, c)}{q(c|x)} \right] \\
&= \sum_{x \in X} \sum_{w_r \in \mathcal{R}_x} \mathbb{E}_{q(C|x)} [\log p(w_r|c)p(c)] \\
&\quad - \sum_{x \in X} \sum_{w_r \in \mathcal{R}_x} \mathbb{E}_{q(C|x)} [\log q(c|x)], \quad (3)
\end{aligned}$$

where x is a document, X is a set of documents, \mathcal{R}_x is the set of signal words of document x , w_r is a signal word, C is a discrete random variable representing the category of a document, c is a possible value of variable C . For example, c can be *science* or *business*.

There are three probabilities in the Eq. (3). $q(c|x)$ is the document classifier which is our ultimate goal. $p(w_r|c)$ is the word distribution learner which estimates the probability distribution of all signal words given a possible value c . We use neural networks to parameterize $p(w_r|c)$ and $q(c|x)$.

$p(C)$ is a prior probability distribution. Since there are no annotated data available, we cannot estimate $p(C)$. Hence we assume it is a uniform distribution, and $p(c)$ becomes a constant.

3.2.1 Word Distribution Learner

The word distribution learner aims to estimate the probability of a signal word w_r given a possible value of category c . It is defined as follows,

$$p(w_r|c) = \frac{\exp(\mathbf{v}_c^T \mathbf{w}_r)}{\sum_{w_{r'}} \exp(\mathbf{v}_c^T \mathbf{w}_{r'})}, \quad (4)$$

where \mathbf{v}_c is a trainable vector associated with c and \mathbf{w}_r is the trainable word embedding of signal word w_r . The intuition is that if a word (e.g., “scientist”) appears frequently under the *science* category, the corresponding inner-product value is high, otherwise it is low.

Eq. (4) requires the summation over all signal words. Since the size of the word vocabulary can be large, we use the negative sampling technique (Mikolov et al., 2013) to approximate Eq. (4). Specifically, we approximate $\log p(w_r|c)$ as follows,

$$\log \sigma(\mathbf{v}_c^T \mathbf{w}_r) + \sum_{w'_r \in \mathcal{N}} \log (1 - \sigma(\mathbf{v}_c^T \mathbf{w}'_r)), \quad (5)$$

where w'_r is a negative sample in the vocabulary, \mathcal{N} is the set of negative samples and $\sigma(\cdot)$ is the sigmoid function.

The objective function with an approximated word distribution learner is defined as follows,

$$\begin{aligned}
\mathcal{L} &= \sum_{x \in X} \sum_{w_r \in \mathcal{R}_x} \mathbb{E}_{q(C|x)} [\log \sigma(\mathbf{v}_c^T \mathbf{w}_r) \\
&\quad + \sum_{w'_r \in \mathcal{N}} \log (1 - \sigma(\mathbf{v}_c^T \mathbf{w}'_r)) + \log p(c)] \\
&\quad - \mathbb{E}_{q(C|x)} [\log q(c|x)]. \quad (6)
\end{aligned}$$

3.2.2 Document Classifier

Most existing deep neural models (DNN) can be used to parameterize $q(C|x)$. As long as the input of DNNs is a document, and the output is a probability distribution of category C . Since models which involve latent variables are difficult to optimize, we give a good initialization of the document classifier. We pre-train the document classifier using pseudo labeled data to initialize it.

Table 2: Statistics and label surface names in AGNews, 20Newsgroup, and UCINews.

Datasets	# Train	# Dev	# Test	# Class	Label Surface Names
AGNews	108,000	12,000	7,600	4	politics, sports, business, technology
20Newsgroup	14,609	1,825	1,825	6	computer graphics, sports car, science electronics encryption health aerospace, politics gun homosexuality, religion atheist christianity, sale
UCINews	26,008	2,560	27,556	4	entertainment, technology, business, health

Table 3: Vocabulary size of signal words that are generated from an MLM and that are extracted from the document (Doc) after removing non-discriminated words.

Dataset	MLM	Doc
AGNews	724	584
20Newsgroup	1,037	413
UCINews	584	442

4 Experiments

In this section, we show the empirical performance of our method on the text classification task.

4.1 Datasets

We evaluate all methods on three datasets.

(1) **AGNews** consists of news articles. It is constructed by (Zhang et al., 2015), which has been gathered from more than 2000 news sources in more than one year of activity.

(2) **20Newsgroup** comprises around 18,000 posts. It is originally collected by (Lang, 1995). We perform text classification on coarse-grained topics. It is an unbalanced dataset.

(3) **UCINews** consists of news pages collected from a web aggregator. It is maintained by (Dua and Graff, 2017).

Table 2 provides statistics and label surface names of three datasets. In 20Newsgroups, we expand label surface names by combining fine-grained label surface names under the same coarse-grained category.

Table 3 shows the vocabulary size of signal words that are generated from an MLM and that extracted from the document (Doc) after removing non-discriminated words.

4.2 Compared Methods

Dataless (Chang et al., 2008) is performed based on vector similarity between documents and label surface names using explicit semantic analysis representation. The prediction is the category that yields the maximum cosine similarity.

Label-Word Similarity is performed based on the vector similarity between words generated from an MLM and label surface names using the static representation. The prediction is the category that yields the maximum cosine similarity.

Pseudo-CNN assigns pseudo labels to documents in the training set based on label-word similarity. We train a CNN model using pseudo labeled samples in the training set. More details are provided in section 4.5.

Pseudo-BERT trains BERT (Devlin et al., 2019) BERT-base-uncased using the same pseudo labeled data as Pseudo-CNN. More details are provided in section 4.5.

WeSTClass (Meng et al., 2018) first generates pseudo labels for documents which contain user-provided keywords. It pre-trains a neural network using pseudo samples as the training set and then performs a self-training process.

LOTClass (Meng et al., 2020) constructs a category vocabulary for each class, using a pre-trained LM. The vocabulary contains words that are relevant to the label name. LOTClass fine-tunes an LM via word-level category prediction task, and then performs self-training on unlabeled data to generalize the model.

ConWea (Mekala and Shang, 2020) leverages contextualized representations of word occurrences and seed word information to automatically distinguish multiple senses of the same word. The

Table 4: Micro F1 and macro F1 scores of all methods on AGNews, 20Newsgroup, and UCINews.

Methods \ Datasets	AGNews		20Newsgroup		UCINews	
	Micro	Macro	Micro	Macro	Micro	Macro
Dataless (Chang et al., 2008)	0.6855	0.6844	0.5000	0.4700	0.6248	0.6253
Label-Word Similarity	0.7917	0.7884	0.7310	0.6390	0.6447	0.6390
Pseudo-CNN	0.8265	0.8237	0.7973	0.6825	0.7598	0.7632
Pseudo-BERT	0.8249	0.8219	0.8153	0.6896	0.7824	0.7820
WeSTClass (Meng et al., 2018)	0.8279	0.8268	0.5300	0.4300	0.6983	0.6999
LOTClass (Meng et al., 2020)	0.8659	0.8656	0.6121	0.5586	0.7320	0.7236
ConWea (Mekala and Shang, 2020)	0.7443	0.7401	0.6200	0.5700	0.3293	0.3269
X-Class (Wang et al., 2021)	0.8574	0.8566	0.6515	0.6316	0.6885	0.6962
WDDC-MLM	0.8826	0.8825	0.8121	0.6882	0.8150	0.8134
WDDC-Doc	0.8668	0.8657	0.8570	0.8250	0.7814	0.7772
CNN (Kim, 2014)	0.9025	0.9025	0.9397	0.9310	0.9002	0.8998
BERT (Devlin et al., 2019)	0.9305	0.9306	0.9660	0.9569	0.9313	0.9315

contextualized corpus is used to train the classifier and expand seed words iteratively.

X-Class (Wang et al., 2021) leverages BERT representations to generate class-oriented document presentations, then generates document-class pairs by clustering, and then fed pairs to a supervised model to train a text classifier.

CNN(Kim, 2014) trains a text CNN using annotated training data in a supervised manner. It is an upper bound of weakly supervised methods.

BERT fine-tunes BERT BERT-base-uncased (Devlin et al., 2019) using annotated training data. It is an upper bound of weakly supervised methods.

WDDC We use a text CNN(Kim, 2014) as the document classifier. Instead of randomly initializing CNN, we pre-train CNN using Pseudo-CNN. **WDDC-MLM** uses the supervision signals from an MLM while **WDDC-Doc** uses the supervision signals from the document itself.

4.3 Result Analysis

Table 4 shows that our method outperforms weakly supervised baselines by 2%, 4%, and 3% in AGNews, 20Newsgroup, and UCINews, respectively. The gaps between the upper bound CNN and our method are 2%, 8%, and 8% in AGNews, 20Newsgroup, and UCINews, respectively. There are still large performance gaps on 20Newsgroup and UCINews.

Label-Word Similarity and Dataless both use vector similarity for prediction. Label-Word Similarity consistently outperforms Dataless, which

shows that words generated from an MLM are useful compared with documents. The performance of Pseudo-BERT is comparable with WeSTClass in AGNews and better than any other baselines in 20Newsgroup and UCINews, which also shows the effectiveness of our pseudo label generation technique. In 20Newsgroup, Macro F1 scores are lower than Micro F1 scores in Pseudo-CNN, Pseudo-BERT, and WDDC-MLM methods. We found that the number of pseudo labeled data of *sale* category is much lower than other categories. So CNN does not have enough pseudo labeled data to learn the *sale* category. The F1 score of *sale* category is lower.

In AGNews and UCINews, WDDC-MLM outperforms WDDC-Doc by 2% and 3%, respectively, which shows that signal words from an MLM are more useful than extracted words from a document. But in 20Newsgroup, WDDC-Doc outperforms WDDC-MLM by 4%. The possible reason is that some categories in 20Newsgroup are not completely disjoint. According to general knowledge, encryption is a field of computer, and computer is a field of science. But in 20Newsgroup (refer to Table 2), science and encryption belong to one class, and computer belongs to another class. MLMs can capture general knowledge from training corpora such as Wikipedia. When given a document talking about encryption, an MLM probably generates words about encryption as well as computer. In this circumstance, generated words are misleading while extracted words are clean. We have detailed

Table 5: Mean and standard deviation of micro and macro F1 scores on 5 independent runs.

Dataset \ Method	WDDC				Baselines			
	Micro F1		Macro F1		Micro F1		Macro F1	
	Mean	Std	Mean	Std	Mean	Std	Mean	Std
AGNews	0.8826	0.0013	0.8825	0.0013	0.8630	0.0038	0.8626	0.0037
20Newsgroup	0.8570	0.0023	0.8250	0.0033	0.8153	0.0131	0.6896	0.0063
UCINews	0.8150	0.0012	0.8134	0.0014	0.7824	0.0141	0.7820	0.0148

Table 6: Some incorrect predictions in AGNews, 20Newsgroup, and UCINews.

Dataset	Text	Prediction	Ground Truth	Signal Words (MLM)
AGNews	Microsoft and Palmone today announced a partnership that will likely have a negative impact on good technology, a well capitalized startup.	technology	business	windows, microsoft, business, security, technology, linux, privacy
20Newsgroup	For the system, or ‘family’, key would appear to be cryptographically useless. ... The same key is used for both encryption and decryption.	computer	science	software, virus, linux, encryption , ibm , nsa
UCINews	Paraplegic teenager to kick off World Cup thanks to robot suit.	entertainment	health	football, sport, soccer cricket, tennis

analysis in section 4.4.

Table 5 shows mean and standard deviation of micro and macro F1 scores of WDDC and best baselines on 5 independent runs. We also conducted t-tests, and p-values are all less than 0.001. We concluded that our method outperforms baselines significantly. Baselines refers to LOTClass, Pseudo-BERT, and Pseudo-BERT on AGNews, 20Newsgroup, and UCINews respectively.

4.4 Case Study

4.4.1 Analysis of Incorrect Predictions

Table 6 shows some incorrect predictions. In the first example, some words in the original document such as “partnership” and “startup” indicate *business* while other words such as “Microsoft” and “technology” indicate *technology*. Signal words generated from an MLM are all related to *technology*. In AGNews dataset, there are a number of samples talking about the stock price of technology companies or cooperation between technology companies. An MLM inclines to focus on either *technology* or *business* and ignore the other one. Although the extraction method can cover all words,

the model is likely to be confused when signal words are related to two categories. In the second example, an MLM generates words related to *encryption* as well as *computer*. Generated words make sense because according to general knowledge, *encryption* is related to *computer*. Unfortunately, most of the signal words from an MLM are related to *computer* except one word “encryption.” WDDC-MLM is likely to predict it as *computer*. Signal words extracted from the document are “encryption” and “key”, which are more likely to guide the model to predict the correct category. In the third example, an MLM generates words that are all about sports because the term “World Cup” appears in the original document. The modifier “paraplegic” plays an important role in identifying the true category. Both generation and extraction methods fail to capture that.

4.4.2 Analysis of Word Distribution Learner

The word distribution learner aims to estimate the probability of a signal word w_r given a possible value of category c , i.e., $p(w_r|c)$. A good word distribution learner should assign a high probability

Table 7: Top 15 signal words that have large inner product values with different latent variable vectors respectively on AGNews dataset. Signal words are generated by an MLM.

Label	Signal Words
Politics	iraq, syria, haiti, israel, murder, baghdad, suicide, torture, war, islam, iran, terrorist, afghanistan, religion, terrorism
Sports	injury, racing, baseball, soccer, boxing, player, relegation, cricket, quarterback, england, basketball, doping, football, golf, tennis
Business	profit, market, finance, agriculture, bankruptcy, energy, money, growth, price, insurance, recession, airline, oil, risk, inflation
Technology	ipod, genetics, encryption, microsoft, internet, hacking, virus, biotechnology, science, copyright, itunes, nasa, evolution, space, astronomy

to category-indicated words, so that by maximizing Eq. (3), a large value of $p(w_r|c)$ leads to a large value of $q(c|x)$, which means if a document contains indicative words to category c , it possibly belongs to category c . Table 7 shows top 15 signal words that have large inner product values with different latent variable vectors respectively on AGNews dataset. As shown in Table 7, the selected words are category-indicated. For example, in the *politics* category, all words are about terrorism, war, and places where wars broke out, which are relevant to the *politics* topic. The word distribution learner can be considered as a category-indicated keywords expansion module.

4.5 Implementation

We use the BERT (`bert-base-uncased`) model to obtain supervision signals in AGNews and 20Newsgroup. We use the BERT (`bert-base-cased`) to obtain supervision signals in UCINews which contains many acronyms such as WHO and PTSD. We select top 20 predictions as supervision signals three datasets. To remove non-discriminated words, we set the threshold t to 2 in three datasets.

In the pseudo label generation process, we set the threshold γ to 0.6, 0.75, and 0.55 in AGNews, 20Newsgroup, and UCINews, respectively. Those pseudo labeled training data are used in Pseudo-CNN and Pseudo-BERT. A higher γ may result in more accurate pseudo labels. But we need to balance the size of pseudo labeled data because it will shrink when γ increases.

To train WDDC, in each batch, we randomly select 5 signal words among all signal words of a document. The number of negative samples in

the approximated word distribution learner is set to 10. For Pseudo-CNN, CNN, and WDDC methods, the CNN architectures are the same. Four different filter sizes $\{2, 3, 4, 5\}$ are applied. A max-pooling layer is applied to each convolutional layer, and each convolutional layer has 100 filters. The maximum length of input in the CNN is set to 64, 128, and 64 in AGNews, 20Newsgroup, and UCINews, respectively. The input in the CNN is contextualized embeddings generated by BERT (`bert-base-uncased`).

For WeSTClass, we use a CNN as the document classifier because it empirically outperforms LSTM in WeSTClass. The CNN architecture we used here is the same as the one described in their paper. We try our best to find good keywords and tune hyper-parameters for WeSTClass and LOT-Class. For all methods, we tune hyper-parameters on development sets.

5 Conclusion

To solve the weakly supervised classification task, we propose to query a masked language model with cloze style prompts to obtain supervision signals. We design a prompt which combines the document itself and “this article is talking about [MASK].” The predictions for the “[MASK]” token are considered as supervision signals because they summarize the content of documents. We propose a latent variable model (WDDC) to learn word distributions given pre-defined categories and a neural document classifier simultaneously without using any annotated data. Evaluation on three datasets shows that our method can outperform weakly supervised learning baselines.

Acknowledgements

The authors of this paper were partially supported by the NSFC Fund (U20B2053) from the NSFC of China, the RIF (R6020-19 and R6021-20) and the GRF (16211520) from RGC of Hong Kong, the MHKJFS (MHP/001/19) from ITC of Hong Kong and the National Key R&D Program of China (2019YFE0198200) with special thanks to Hong Kong Mediation and Arbitration Centre (HKMAAC) and California University, School of Business Law & Technology (CUSBLT), and the Jiangsu Province Science and Technology Collaboration Fund (BZ2021065). We also thank the anonymous reviewers for their valuable comments and suggestions that help improve the quality of this manuscript.

References

- Cornelia Caragea, Nathan J. McNeese, Anuj R. Jaiswal, Greg Traylor, Hyun-Woo Kim, Prasenjit Mitra, Dinghao Wu, Andrea H. Tapia, C. Lee Giles, Bernard J. Jansen, and John Yen. 2011. Classifying text messages for the haiti earthquake. In *Proceedings of ISCRAM*.
- Mingwei Chang, Lev Arie Ratinov, Dan Roth, and Vivek Srikumar. 2008. Importance of semantic representation: dataless classification. In *Proceedings of AACL*, pages 830–835.
- Wenhu Chen, Wenhan Xiong, Xifeng Yan, and William Wang. 2018. Variational knowledge graph reasoning. In *Proceedings of NAACL-HLT*, pages 1823–1832.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT*, pages 4171–4186.
- Dheeru Dua and Casey Graff. 2017. [UCI machine learning repository](#).
- Evgeniy Gabrilovich, Shaul Markovitch, et al. 2007. Computing semantic relatedness using wikipedia-based explicit semantic analysis. In *Proceedings of IJCAI*, pages 1606–1611.
- Johan Ludwig William Valdemar Jensen et al. 1906. Sur les fonctions convexes et les inégalités entre les valeurs moyennes. *Acta mathematica*, 30:175–193.
- Rie Johnson and Tong Zhang. 2017. Deep pyramid convolutional neural networks for text categorization. In *Proceedings of ACL*, pages 562–570.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. In *Proceedings of EMNLP*, pages 1746–1751.
- Diederik P Kingma and Max Welling. 2014. Auto-encoding variational bayes. In *Proceedings of ICLR*.
- Toutanova Kristina, Klein Dan, Manning Christopher, and Yoram Singer. 2003. Feature-rich part-of-speech tagging with a cyclic dependency network. In *Proceedings of NAACL-HLT*, pages 252–259.
- Ken Lang. 1995. Newsweeder: Learning to filter news. In *Proceedings of ICML*, pages 331–339.
- Yan Liang, Xin Liu, Jianwen Zhang, and Yangqiu Song. 2019. Relation discovery with out-of-relation knowledge base as supervision. In *Proceedings of NAACL-HLT*, pages 3280–3290.
- Han Liu, Xiaotong Zhang, Lu Fan, Xuandi Fu, Qimai Li, Xiao-Ming Wu, and Albert Y. S. Lam. 2019. Reconstructing capsule networks for zero-shot intent classification. In *Proceedings of EMNLP*, pages 4798–4808.
- Diego Marcheggiani and Ivan Titov. 2016. Discrete-state variational autoencoders for joint discovery and factorization of relations. *Transactions of the Association for Computational Linguistics*, 4:231–244.
- Dheeraj Mekala and Jingbo Shang. 2020. Contextualized weak supervision for text classification. In *Proceedings of ACL*, pages 323–333.
- Yu Meng, Jiaming Shen, Chao Zhang, , and Jiawei Han. 2018. Weakly-supervised neural text classification. In *Proceedings of CIKM*, pages 983–992.
- Yu Meng, Jiaming Shen, Chao Zhang, , and Jiawei Han. 2019. Weakly-supervised hierarchical text classification. In *Proceedings of AACL*, pages 6826–6833.
- Yu Meng, Yunyi Zhang, Jiaxin Huang, Chenyan Xiong, Heng Ji, Chao Zhang, and Jiawei Han. 2020. Text classification using label names only: A language model self-training approach. In *Proceedings of EMNLP*, pages 9006–9017.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Proceedings of NeurIPS*, pages 3111–3119.
- Fabio Petroni, Patrick Lewis, Aleksandra Piktus, Tim Rocktäschel, Yuxiang Wu, Alexander H Miller, and Sebastian Riedel. 2020. How context affects language models’ factual predictions. In *Proceedings of AKBC*.
- Fabio Petroni, Tim Rocktäschel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, Alexander H Miller, and Sebastian Riedel. 2019. Language models as knowledge bases? In *Proceedings of EMNLP*, pages 2463–2473.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI blog*.

- Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. 2014. Stochastic backpropagation and approximate inference in deep generative models. In *Proceedings of ICML*, pages 1278–1286.
- Anthony Rios and Ramakanth Kavuluru. 2018. Few-shot and zero-shot multi-label learning for structured label spaces. In *Proceedings of EMNLP*, pages 3132–3142.
- Herbert L. Roitblat, Anne Kershaw, and Patrick Oot. 2010. Document categorization in legal electronic discovery: computer classification vs. manual review. *Journal of the Association for Information Science and Technology*, 61(1):70–80.
- Timo Schick and Hinrich Schütze. 2021. Exploiting cloze questions for few shot text classification and natural language inference. In *Proceedings of EACL*, pages 255–269.
- Timo Schick and Hinrich Schütze. 2021. It’s not just size that matters: Small language models are also few-shot learners. In *Proceedings of NAACL-HLT*, pages 2339–2352.
- Taylor Shin, Yasaman Razeghi, Robert L Logan IV, Eric Wallace, and Sameer Singh. 2020. Autoprompt: Eliciting knowledge from language models with automatically generated prompts. In *Proceedings of EMNLP*, pages 4222–4235.
- Yangqiu Song and Dan Roth. 2014. On dataless hierarchical text classification. In *Proceedings of AACL*, pages 1579–1585.
- Simon Šuster, Ivan Titov, and Gertjan van Noord. 2016. Bilingual learning of multi-sense embeddings with discrete autoencoders. In *Proceedings of NAACL-HLT*, pages 1346–1356.
- Ivan Titov and Ehsan Khoddam. 2015. Unsupervised induction of semantic roles within a reconstruction-error minimization framework. In *Proceedings of NAACL-HLT*, pages 1–10.
- Zihan Wang, Dheeraj Mekala, and Jingbo Shang. 2021. X-class: Text classification with extremely weak supervision. In *Proceedings of NAACL-HLT*, pages 3043–3053.
- Congying Xia, Chenwei Zhang, Xiaohui Yan, Yi Chang, and Philip S. Yu. 2018. Zero-shot user intent detection via capsule neural networks. In *Proceedings of EMNLP*, pages 3090–3099.
- Ziqian Zeng, Wenxuan Zhou, Xin Liu, and Yangqiu Song. 2019. A variational approach to weakly supervised document-level multi-aspect sentiment classification. In *Proceedings of NAACL-HLT*, pages 386–396.
- Jingqing Zhang, Piyawat Lertvittayakumjorn, and Yike Guo. 2019. Integrating semantic knowledge to tackle zero-shot text classification. In *Proceedings of NAACL-HLT*, pages 1031–1040.
- Rongzhi Zhang, Yue Yu, Pranav Shetty, Le Song, and Chao Zhang. 2022. Prboost: Prompt-based rule discovery and boosting for interactive weakly-supervised learning. In *Proceedings of ACL*.
- Xiang Zhang, Junbo Jake Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. In *Proceedings of NeurIPS*, pages 649–657.
- Yuyu Zhang, Hanjun Dai, Zornitsa Kozareva, Alexander J Smola, and Le Song. 2018. Variational reasoning for question answering with knowledge graph. In *Proceedings of AAAI*, pages 6069–6076.
- Peng Zhou, Zhenyu Qi, Suncong Zheng, Jiaming Xu, Hongyun Bao, and Bo Xu. 2016. Text classification improved by integrating bidirectional lstm with two-dimensional max pooling. In *Proceedings of COLING*, pages 3485–3495.