

FastClass: A Time-Efficient Approach to Weakly-Supervised Text Classification

Tingyu Xia^{1,3}, Yue Wang², Yuan Tian^{1,3,*}, Yi Chang^{1,3,4,*}

¹School of Artificial Intelligence, Jilin University

²School of Information and Library Science, University of North Carolina at Chapel Hill

³Key Laboratory of Symbolic Computation and Knowledge Engineering, Jilin University

⁴International Center of Future Science, Jilin University

xiaty21@mails.jlu.edu.cn, wangyue@unc.edu

yuantian@jlu.edu.cn, yichang@jlu.edu.cn

Abstract

Weakly-supervised text classification aims to train a classifier using only class descriptions and unlabeled data. Recent research shows that keyword-driven methods can achieve state-of-the-art performance on various tasks. However, these methods not only rely on carefully-crafted class descriptions to obtain class-specific keywords but also require substantial amount of unlabeled data and takes a long time to train. This paper proposes **FastClass**, an efficient weakly-supervised classification approach. It uses dense text representation to retrieve class-relevant documents from external unlabeled corpus and selects an optimal subset to train a classifier. Compared to keyword-driven methods, our approach is less reliant on initial class descriptions as it no longer needs to expand each class description into a set of class-specific keywords. Experiments on a wide range of classification tasks show that the proposed approach frequently outperforms keyword-driven models in terms of classification accuracy and often enjoys orders-of-magnitude faster training speed.

1 Introduction

Text classification is one of the most used techniques in mining large-scale unstructured text. When sufficient labeled data are available, supervised classification techniques can achieve excellent performance. However, manually labeling example documents can be time-consuming and labor-intensive, a major burden when applying supervised text classification techniques in practice.

Recently, *weakly-supervised text classification* (Meng et al., 2018, 2019, 2020; Mekala and Shang, 2020; Shen et al., 2021; Wang et al., 2021; Zhang et al., 2021) has been proposed to save labeling efforts. It refers to the ability for a machine learning model to start classifying documents by using only class descriptions and unlabeled data.

Since weakly-supervised text classification methods do not rely on any labeled data, it can greatly reduce the workload of data labeling. Therefore, these methods are desirable in real-world text mining applications.

The current mainstream weakly-supervised text classification methods are keyword-driven (Meng et al., 2018, 2019, 2020; Mekala and Shang, 2020; Shen et al., 2021; Wang et al., 2021; Zhang et al., 2021), where the users need to carefully choose initial class-relevant keywords for each class. Such keywords are expanded into a richer set of keywords, which are then used to assign pseudo-labels on unlabeled data. The critical step of these keyword-driven methods is the choice of initial class-relevant keywords and the subsequent keyword expansion strategy, which determines the quality of pseudo-labeled documents.

However, if the user’s initial choice of class descriptions is ambiguous or brief, then the results of keyword expansion will be negatively impacted. Ideally, the user can provide initial class descriptions that are beneficial to keywords expansion. However, in reality, the quality of class descriptions is not guaranteed. Early research often artificially modifies original class descriptions to improve the quality of the initial keywords (Chang et al., 2008). In sentiment classification tasks, “positive” and “negative” class descriptions have to be replaced by “good” and “bad”, which are easier for keywords expansion (Meng et al., 2020).

Previous weakly-supervised methods also face the common problem of fully supervised methods. That is, model performance depends on the amount of (unlabeled) data available. Widely used classification datasets in research literature often have high-quality class descriptions and sufficient unlabeled data. However, weakly-supervised classification methods are most needed in the early stage of practical text mining tasks, where practitioners are not ready to invest vast amounts of efforts in data

*Corresponding authors.

labeling or computing resources in model training. In this stage, class descriptions can be crude and even unlabeled data can be scarce (e.g., classifying text related to an emerging event).

As we will show in our experiments, the increasingly sophisticated keyword-driven weakly-supervised methods require increasing amount of computing resources and time in model training, and yet the resulting performance benefits can be minimal. This is concerning given the recent recommendation on Green AI (Schwartz et al., 2020) and Efficient NLP (Arase et al., 2021). Recent research found that Transformer-based textual entailment models can provide more competitive performance on dataless classification tasks (Yin et al., 2019; Chu et al., 2020) which is similar to weakly-supervised text classification. In principle, such models only need to be trained once and can be applied to any classification tasks. However, it is not only difficult to adapt such models to a specific corpus, but also inefficient to run them at prediction time. That is, one has to run the entailment model k times to classify one document into k categories. The prediction speed slows down as more categories (larger k) are considered in a task.

In this paper, we propose **FastClass**, a time-efficient weakly-supervised text classification model which constructs a classification model by selecting an optimal subset of retrieved documents to solve the above problems. **FastClass** has three steps: (1) extracting seed unlabeled documents with high semantic similarity to the original class descriptions; (2) using seed documents as queries to obtain pseudo-labeled documents from external corpus. (3) selecting an optimal subset of pseudo-labels using the maximum entropy principle.

Our main contributions are as follows:

- We propose an efficient weakly-supervised text classification method that selects a document subset returned by dense retrieval models as pseudo-labels for classifier training.
- Compared to keyword-driven weakly-supervised methods, our method is less sensitive to class descriptions and does not heavily rely on high-quality category descriptions.
- Extensive empirical experiments show that our method has higher accuracy and faster training speed in most cases, and the scale of

task data does not have a significant impact on training time and model performance.

2 Related Work

We discuss related work from three perspectives: zero-shot text classification, dataless text classification, and weakly-supervised text classification.

2.1 Zero-Shot Text Classification

Zero-shot text classification divides classes into seen classes with annotated data and unseen classes without any labeled data. Since no labeled data is available for unseen classes during training, the general idea of zero-shot learning is to transfer knowledge from seen classes to unseen classes.

The mainstream zero-shot text classification method is to exploit semantic knowledge to generally infer the features of unseen classes using patterns learned from seen classes. Three main types of semantic knowledge have been employed in general zero-shot scenarios, including semantic attributes and properties of classes (Liu et al., 2019a; Xia et al., 2018; Pushp and Srivastava, 2017), correlations among classes (Rios and Kavuluru, 2018; Zhang et al., 2019), semantic word embeddings which capture implicit relationships between classes and documents (Nam et al., 2016). Besides, Ye et al. explored reinforcement learning framework to tackle zero-shot task.

2.2 Dataless Text Classification

Dataless text classification (Chang et al., 2008) aims to classify text using a given set of class descriptions and no labeled data for training a model. These methods have two broad categories: classification-based (Song and Roth, 2014; Yin et al., 2019; Chu et al., 2020) and clustering-based (Li et al., 2018; Li and Yang, 2018).

In previous works, dataless text classification also has many slightly different setups. For example, in zero-shot text classification, Yin et al. proposed “label-fully-unseen” setting which directly computes document-label relatedness with a sentence-pair BERT model. The model is trained with large-scale texts naturally tagged with category information, such as Wikipedia. NAT-CAT (Chu et al., 2020) combines various publicly available online corpora that come with natural categories, and trains a BERT (Devlin et al., 2018) or RoBERTa (Liu et al., 2019b) model to discriminate correct versus incorrect categories for a given

document. These methods create pseudo-labeled data from external resources to train a universal textual entailment model that can be applied to a wide spectrum of classification tasks.

2.3 Weakly-supervised Text Classification

It is easy to confuse “dataless text classification” with “zero-shot text classification” (Wang et al., 2019; Ye et al., 2020) and “weakly-supervised text classification” (Meng et al., 2019, 2020). Zero-shot text classification may still provide labeled data for part of the categories (label-partially-seen (Yin et al., 2019)), while dataless text classification can be applied to any text classification task with a given set of label descriptions by building a general text classifier. Weakly-supervised text classification assumes a large amount of unlabeled data and label descriptions are available for training.

Most previous works on weakly-supervised text classification are keyword-driven methods. These methods generate pseudo-labeled data by counting class-relevant keywords in unlabeled documents and then train a supervised model. Among them, WeSTClass (Meng et al., 2018) utilizes a self-training module which bootstraps on unlabeled data for model refinement. WeSHClass (Meng et al., 2019) extends WeSTClass to hierarchical labels. LOTClass (Meng et al., 2020) uses label names as initial keywords and augments the keywords with BERT’s MLM module. ConWea (Mekala and Shang, 2020) disambiguates polysemy keywords based on contextualized corpus. X-Class (Wang et al., 2021) constructs pseudo-labeled document by estimating class-oriented document representation and document clustering. ClassKG (Zhang et al., 2021) improves the quality of pseudo-labels by applying GNN to a keyword graph to exploit keyword correlation.

3 Proposed Methods

In this section, we describe our proposed method **FastClass** for weakly-supervised text classification. We formulate the problem as follows. We are given a set of class descriptions $D = \{d_1, \dots, d_j, \dots, d_k\}$, each is a piece of short text (one or more words) describing a semantic class j in the label space $Y = \{1, \dots, k\}$. We are given a set of unlabeled documents X in the task domain. As a natural scenario in practice, we also have access to vast amounts of external unlabeled documents U , $|U| \gg |X|$. These external docu-

ments may come from Wikipedia, news corpora, and online social media, which may or may not share the same domain as the classification task in question. Our goal is to correctly assign label(s) from Y to (a subset of) unlabeled documents in U as pseudo-labeled training data.

At a high level, our proposed method uses class descriptions in D to obtain task-specific documents that are highly similar to class descriptions from X and uses these task-specific documents as queries to retrieve pseudo-labeled documents from external unlabeled data U . Then we select the optimal subset of these pseudo-labeled documents to train a classifier. The general **FastClass** framework is presented in Figure 1. Below we describe our method in detail.

3.1 Dense Text Representation and Indexing

As a preparation step, we use a sentence representation model to convert all texts (class descriptions, task-specific unlabeled documents, and external unlabeled documents) into dense vectors in a semantic space. In principle, any dense text representation techniques can be used. We choose to use SentenceBERT (SBERT) (Reimers and Gurevych, 2019) as it is proven to deliver good performance in various sentence-pair modeling and information retrieval tasks (Thakur et al., 2021).

Once these texts are converted into dense vectors, we build approximate nearest neighbor (ANN) indices for task-specific unlabeled documents and external documents to enable fast document retrieval. In principle, any ANN search techniques can be used. We choose to use FAISS (Johnson et al., 2017) for efficient similarity search with cosine similarity as the vector similarity metric. We also tested other metrics such as Euclidean distance but found negligible performance difference.

3.2 Class-Relevant Document Retrieval

The first step of our method is to retrieve a pool of potentially relevant documents for each class, a subset of which are pseudo-labeled in the next step.

We retrieve pseudo-labeled documents from external data with a task-specific focus. The idea is to enrich a class description with task-specific data before retrieving from external documents. For each class j , we first obtain a “seed set” of documents S_j . We use each class description as a search query to retrieve documents from task-specific unlabeled data. For class $j \in Y$, we rank documents in the unlabeled data X by their semantic similarity to

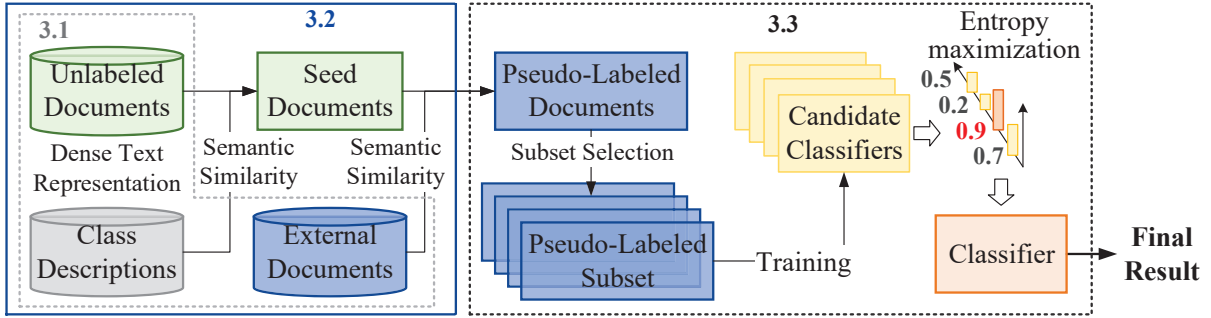


Figure 1: An overview of the **FastClass** approach. We first get dense representations of unlabeled documents, class descriptions and external documents (Section 3.1). Then we obtain task-specific seed documents from unlabeled documents and use them as queries to retrieve pseudo-labeled documents from external documents (Section 3.2). Lastly, we select the optimal subset of these pseudo-labeled documents based on entropy maximization strategy (Section 3.3). Different colors represent different origins of data.

the class description d_j and take the most similar c documents $S_j = \{x_i\}_{i=1}^c$. Here, semantic similarity is computed using the vectors produced in Section 3.1. Then we use S_j to further retrieve external documents by treating each $x \in S_j$ as a query to retrieve its n nearest neighbors $\Gamma(x)$ from external data. However, these documents may be close to a seed document because they share words unrelated to the theme of the class. To filter such noise, we preserve documents that appear in at least two seed documents’ nearest neighborhoods. This gives class-relevant documents for class j : $R_j = \{e | \exists x_1, x_2 \in S_j, e \in \Gamma(x_1) \wedge e \in \Gamma(x_2)\}$. The hope is that R_j contains external documents that are semantically relevant and stylistically similar to task-specific unlabeled data.

3.3 Pseudo-Labeled Subset Selection

Subset diversification. We select a subset $L_j \subset R_j$ of size m to be pseudo-labeled as class j . The motivation is that documents retrieved from external data sources may contain (near-)duplicates. For example, many news outlets may cover the same story. Duplicated documents may lead to overfitting as they give too much emphasis on a few documents and reduce the overall diversity of pseudo-labeled training data. Indeed, previous works have shown that diverse training data improves learning performance (Wei et al., 2015). Here we apply facility location function to quantify the diversity of a subset. The facility location function of any subset $L_j \subset R_j$ is defined as

$$g(L_j) = \sum_{x \in R_j} \max_{e \in L_j} s(x, e). \quad (1)$$

Here $s(\cdot, \cdot)$ is the cosine similarity between two dense document vectors. Intuitively, $g(L_j)$ computes the total cost for every element $x \in R_j$ to be “covered” by the most similar element $e \in L_j$. In our context, this translates into how well the subset L_j preserves the content of the larger set R_j . Although finding the optimal subset L_j that maximizes the submodular function $g(L_j)$ is NP-hard, a greedy algorithm gives an approximately optimal solution (Nemhauser et al., 1978). The algorithm sequentially adds the next element x to L_j with the maximum marginal gain $g(L_j \cup \{x\}) - g(L_j)$, until L_j reaches the desired size m .

A challenging problem remains: how many documents to retrieve and assign pseudo-labels (namely, how to set n and m)? More generally, what is the optimal subset of retrieved documents that, if pseudo-labeled, will train a good classifier? Note that we cannot tune subset selection procedures on labeled data as such data is unavailable in a weakly-supervised setting! To address this problem, we propose a novel *unsupervised* subset selection procedure as follows.

Entropy maximization. We now determine the subset selection parameters $\theta = \{n, m\}$. θ determines the pseudo-labeled set L_j for class j , which determines the full pseudo-labeled set $\cup_{j=1}^k L_j$, which in turn trains a classifier $f : X \rightarrow Y$. Below we use f_θ to emphasize that f depends on θ . Once trained, f_θ induces a distribution over the label space Y when applied to the task-specific unlabeled data X : $\forall y \in Y$,

$$p(y|X, f_\theta) = \frac{\sum_{x \in X} \mathbf{1}\{f_\theta(x) = y\}}{|X|}. \quad (2)$$

According to the *maximum entropy* principle




Figure 2: Handling the *Other* class.

(Jaynes, 1957), the distribution with maximum entropy shall be preferred since we have *no* labeled data as evidence to prefer any other distribution. Following this principle, we seek for θ that maximizes the f_θ -induced classification entropy:

$$H(\theta) = \sum_{y \in Y} -p(y|X, f_\theta) \log p(y|X, f_\theta). \quad (3)$$

Empirically, $H(\theta)$ correlates well (but not perfectly) with true performance of f_θ on labeled data even though it is an unsupervised metric (Appendix C), a phenomenon first observed in (Baram et al., 2004). As $H(\theta)$ is non-differentiable with respect to θ , we resort to grid search. It is sufficient to use a coarse grid to find sensible θ values (Section 4.3).

3.4 Handling the *Other* Class

In some classification tasks, we have clearly defined categories and an *Other* category, such as an “other topic” category in topic classification or a “no emotion” category in emotion classification. We call clearly defined (non-*Other*) categories *named classes*. Using “other topic” or “no emotion” literally as the search query to retrieve pseudo-labeled documents is problematic because the *Other* class is to be interpreted with respect to named classes.

The general idea is to pseudo-label documents that are far from any named class as the *Other* class. Without loss of generality, let the named classes be numbered from 1 to $k - 1$ and the *Other* class be class k .

We first select *Other* documents L_k from task-specific unlabeled data $O = X \setminus \cup_{j=1}^{k-1} L_j$. Our goal is to find a subset $L_k \subset O$ with size c that is farthest from the descriptions of all named classes $D \setminus \{d_k\}$. We seek for the subset that *minimizes* the following function:

$$h(L_k) = \sum_{x \in L_k} \max_{e \in D \setminus \{d_k\}} s(x, e). \quad (4)$$

This function is modular and can be efficiently minimized by selecting c documents that have smallest $\max_{j=1}^{k-1} s(d_j, x)$ values from O (Figure 2 left). This turns *Other* into another named class. We then retrieve and select pseudo-labels from external data using the same procedure described in Sections 3.2 and 3.3 (Figure 2 right).

4 Experimental Evaluation

4.1 External Document Repository

To cover various task domains, we combine five large-scale datasets as the external document repository. These datasets are freely available and frequently used in previous works as external resources. We keep these documents short (e.g. titles) as SBERT is well-trained on sentence pairs. We build a single index for all the external documents.

Microsoft News Dataset (MIND) (Wu et al., 2020) is collected from anonymized behavior logs of Microsoft News website. **Multi-Domain Sentiment Dataset (MDS)** (Blitzer et al., 2007) contains product reviews for many product categories in Amazon. **Wikipedia-500K** (Bhatia et al., 2016) has over a million curator-generated category labels and each article often has more than one relevant labels. We select the first sentence of each article. **RealNews** (Zellers et al., 2019) is a large news corpus from Common Crawl. We randomly sample 2M titles from these 32M news. **S2ORC** (Lo et al., 2020) is a general corpus of scientific literature. We randomly select 100k papers from all 20 research fields and extract their titles. The statistics of datasets are shown in Table 1.

4.2 Evaluation Datasets

We choose 9 text classification tasks in our experiments. Note that we only use the original class descriptions (see Appendix A).

Yahoo (Zhang et al., 2015) consists of 10 categories of questions in online forums. **20News-**

Dataset	#Docs	#Words/doc
MIND	285,875	10.7
MDSB	821,249	137.5
Wikipedia	1,779,881	22.9
RealNews	2,000,000	9.6
S2ORC	2,000,000	10.9

Table 1: Statistics of external document datasets.

group (Lang, 1995) is a collection of 20 topic news-group documents. **AGnews** (Zhang et al., 2015) contains 4 topical categories of news titles. **DBPedia** (Lehmann et al., 2015) contains titles, descriptions, and associated categories from DBpedia. **Yelp** (Zhang et al., 2015) is for sentiment analysis in Yelp reviews. **Emotion** (Oberländer and Klinger, 2018) was constructed by combining multiple public datasets where documents have emotion labels. **Amazon** (Zhang et al., 2015) is a binary sentiment classification dataset. **SST** (Socher et al., 2013) is a corpus extracted from movie reviews. **Situation** (Zhang et al., 2015) is a event-type classification dataset originally designed for low-resource situation detection. The statistics of datasets are shown in Table 2.

Dataset	#Classes	#train	#test
Yahoo	10	1300k	100k
AGnews	4	120k	7,600
20News	20	11,015	7,318
DBPedia	14	560k	70k
Yelp	2	560k	38k
Emotion	10	34,667	16k
Amazon	2	3600k	400k
SST-B	2	67,349	1,821
Situation	12	4,921	3,495

Table 2: Statistics of evaluation datasets.

4.3 Compared Methods

We include two state-of-the-art methods for data-less text classification and three weakly-supervised text classification methods:

Dataless text classification. *Label-fully-unseen* **0SHOT-TC** (Yin et al., 2019) pushes “zero-shot learning” to the extreme – no annotated data for any labels. It aims to classify documents without seeing any task-specific training data. **NATCAT** (Chu et al., 2020) proposed to use large-scale, naturally

annotated data to train robust entailment-based text classification models. These two methods both use readily available resources to train textual entailment models that can robustly handle a wide range of text classification tasks.

Weakly-supervised text classification. **LOT-Class** (Meng et al., 2020) utilizes pre-trained LM to collect category-indicative words and generalizes the model via document level self-training on abundant unlabeled data. **X-Class** (Wang et al., 2021) estimates class-oriented document representations based on pre-trained language models from a representation learning perspective, and then selects high-confidence clustered examples to form a pseudo-training set. **ClassKG** (Zhang et al., 2021) designs a new pretext task based on the keyword graph to learn better representations of keyword subgraphs, with which the accuracy of pseudo-label generation is improved, and thus improves classification performance.

4.4 Experimental Settings

Below, we summarize the implementation details that are key for reproducing results.

We use “paraphrase-MiniLM-L6-v2” as the base model for SBERT to obtain the sentence embeddings and the dimension of embedding vectors is 384. FAISS is used to retrieve external documents which works with inner product to compute cosine similarity. The number of clusters is set to 512 and 3 clusters are explored at search time. We implemented facility location subset selection using the Apricot library (Schreiber et al., 2020), which provides cosine as a similarity measure and a lazy greedy optimizer as a solver.

In our pilot study, we experimented with three approaches (Annoy, HNSWlib and FAISS) to building the approximate nearest neighbor index in the initial experimental phase. There was no substantial difference in their performance, so we chose FAISS which enjoys the fastest speed.

When retrieving seed documents, we set $c = .1 \times |X|/k$ to avoid too few pseudo-labels or too many inaccurate pseudo-labels per class. As a weakly supervised method, FastClass can directly operate on unlabeled documents without distinguishing training vs. test sets provided by a task. Therefore, we simply retrieve seed documents from unlabeled (i.e., test) documents to select an appropriate number of seeds, which adapt to the data size of the documents to be classified.

Method	Yahoo	AGnews	20News	DBPedia	Yelp	Emotion	Amazon	SST	Situation
OSHOT-TC	34.8	53.8	22.2	53.8	73.4	<u>21.7</u>	76.0	<u>71.7</u>	16.2
NATCAT	48.6	74.9	44.8	85.3	50.1	10.7	50.8	50.5	27.4
LOTClass	53.3	83.1	/	90.9	/	/	51.2	49.6	/
X-Class	47.7	<u>85.7</u>	39.2	91.5	63.2	13.1	84.2	48.3	41.1
ClassKG	<u>60.8</u>	87.9	<u>49.4</u>	98.3	92.5	/	<u>88.0</u>	67.0	/
FastClass	61.9	85.1	55.4	<u>92.4</u>	<u>87.9</u>	28.0	89.1	84.5	<u>39.2</u>
SUPERVISED									
RoBERTa	75.1	95.4	73.5	99.3	97.5	37.8	97.4	95.8	58.4

Table 3: Weakly-supervised text classification performance on nine datasets (%). The metrics are label-weighted F1 for Emotion and Situation and accuracy for other tasks. Mentioned that “/” indicates the current model can not obtain classification result of the current dataset. The best performance in each column is bold and second-best performance is underlined.

Method	Yahoo	AGnews	20News	DBPedia	Yelp	Emotion	Amazon	SST	Situation
LOTClass	15h↓	29m↓	/	5.5h↓	/	/	15h↓	13.5m↓	/
X-Class	32.5h↓	75m↑	0.5h↓	8.3h↓	14.3h↓	24m↓	62h↓	35m↓	8m↑
ClassKG	41.7h↓	90m↑	8.5h↓	13.7h↑	26.5h↑	/	133.5h↓	54m↓	/
FastClass	1.7h	22m	1h	1.5h	40m	44m	3.5h	12m	33m

Table 4: Comparison of training time for weakly supervised methods. The arrow on the right shows the performance of the current model compared to FastClass. “h”: hours; “m”: minutes. “/” has the same meaning as in Table 3.

To select pseudo-labeled subsets that have maximum classification entropy, we searched parameters θ on the grids $n = \{100, 200, 300\}$ and $m = \{300, 500, 800\}$. The subset-induced RoBERTa-base (110M parameters)(Liu et al., 2019b) classifier that achieved the maximum entropy was used. The optimizer is AdamW (Loshchilov and Hutter, 2017), learning rate is $2e^{-5}$, training batch size is 32 and the number of training epochs is 4.

The training and evaluation of all methods are performed on a NVIDIA GeForce RTX 3090 GPU with 24GB memory and an Intel Xeon Gold 6330 CPU with 14 cores and 80GB memory.

5 Results

In this section, we evaluate our proposed method and compare it with baseline models for dataless and weakly-supervised text classification. The comparison is not only in terms classification accuracy, but also model efficiency.

5.1 Performance Across Datasets

Table 3 summarizes classification performance of baseline methods and our pseudo-labeling method FastClass combined with RoBERTa classifier. We also show classification performance of FastClass

based on BERT classifier in Appendix Table 8. For dataless methods, we report results of OSHOT-TC based on the official pre-trained model used MNLI as training resource and our pre-trained NATCAT based on RoBERTa using Wikipedia as training resource. For weakly-supervised methods, we report results of our reproduction of LOTClass, X-Class, and ClassKG. All baseline results are obtained by ourselves based on the official codes. Besides, we have stored our implementation as open source code in our Github repository¹.

We followed the metrics used in (Yin et al., 2019) and chose label-weighted F1 as the metrics for unbalanced datasets Emotion and Situation. For other balanced datasets, we used accuracy. Note that our method works on a weakly supervised setting, where the data has only texts but no labels. Our approach to generating suitable pseudo-labels for unlabeled texts is deterministic and does not involve randomness. Performing k -fold cross validation over pseudo-labels does not fit in our approach because classifier training is not the final step, but an inner-loop step of the entropy maximization procedure. So statistical significance tests are not applied

¹<https://github.com/xiatingyu/FastClass>

in this work.

Here we make a special remark on the Situation and Emotion datasets: they both contain the *Other* class. For Situation this category is “out-of-domain” and for Emotion it is “no emotion”. We handled the *Other* classes using the approach in Section 3.4.

To compare all methods more fairly, we use a unified training set, testing set and class descriptions. Since ClassKG can achieve good results after 3 iterations, in order to save time, we set the number of iterations to 3. Note that we use *original* class descriptions (e.g., “positive”, “negative”) without rewriting them (e.g., as “good”, “bad”) (see Appendix A). Therefore, the results of LOTClass, X-Class, and ClassKG are different from those reported in the original papers on tasks Amazon, Yelp, and 20News.

In addition, because LOTClass can not find enough unlabeled documents containing class-relevant keywords as pseudo-label data on some tasks, some results are missing. For the datasets Emotion and Situation, during the self-training process of ClassKG, extreme predictions occurred, resulting in no training data for some categories, and the training process stopped.

Although our weakly supervised method does not reach the best level, it shows good results on the whole, especially for some datasets such as SST e.g. which has a small amount of training data, our method can well solve the problem of sparse data. In addition, our method is not highly dependent on labels, which only needs to use the original labels to achieve good results.

5.2 Efficiency Comparison

Training Efficiency. We compare the training time of different weakly supervised methods in Table 4. Our method FastClass does not achieve the best results on all tasks, but FastClass can significantly save training time. Taking the dataset DBpedia as an example, the experimental results of ClassKG are significantly higher than all other methods, but its training time is more than 9 times that of FastClass. Even for the small-scale dataset SST, ClassKG takes about an hour to complete the training, but it cannot achieve the optimal results. For large-scale datasets, the FastClass method achieves relatively good results without sacrificing huge training time. For small-scale datasets, FastClass makes up for the problem of insufficient training due to sparse training data. Figure 3 shows Fast-

Class can get competitive results using the least amount of time.

Here we need to make a point that we did not consider the computing time for sentence representation in Table 4, because it is a one-time, once-and-for-all process. The sentence embedding process for the entire external document repository took less than an hour, and we provided the Python pickle file of the embedding as part of the data and code release for other researchers to use in our Github repository.

Method	Total Time	Per Document
0SHOT-TC	2162.4s	22ms
NATCAT	1485.8s	15ms
FastClass	306.7s	3ms

Table 5: Total testing time on Yahoo using *label-fully-unseen* 0SHOT-TC, NATCAT and FastClass. All methods used RoBERTa-base model.

Prediction efficiency. Compare with the commonly used entailment model which only need to be trained once to be applied to any task in dataless tasks. A big advantage of classification models over entailment models is the prediction speed. Classification models only need one forward pass to make a prediction for k categories, whereas entailment models need k forward passes. Table 5 compares prediction time of entailment models and FastClass on the Yahoo dataset (100,000 documents). Our method is not only more accurate (Table 3) but also 5-7 times faster.

5.3 Ablation Studies

Seed sets: To study the benefit of using task-specific unlabeled data as “seed sets”, we directly construct a pseudo training set based on the similarity between the class description and the external documents. We call this variant FastClass_{external}. It directly uses the class descriptions as queries to retrieve data from external data. For class j , we retrieve n' most relevant documents from external data with respect to the class description d_j , $R_j = \{x_i\}_{i=1}^{n'}$. Then, we use the retrieved external data for subset selection, which is consistent with FastClass include subset diversification and entropy maximization. The parameter θ is $n' = \{2m, 5m, 10m\}$ and $m = \{300, 500, 800\}$.

Table 6 shows that instead of using task specific data as retrieval queries, directly using class descriptions for retrieval will affect the experimental




Figure 3: Performance and efficiency comparison of weakly-supervised text classification models. Polyline: Accuracy of different models. Histogram: training time for different models.

Method	AGnews	Dbpedia	SST	Yelp
FastClass _{external}	72.7	84.9	80.1	87.9
FastClass	85.1	92.4	84.5	87.9

Table 6: Impact of seed sets on performance (%). The best performance of each column is in bold.

results. Although there is a semantic gap between task data and class descriptions, task data can represent the textual features of specific task, such as text style, which is beneficial for retrieving semantically related and stylistically similar documents.

Facility location function: To study the impact of the facility location subset selection function on the diversity of pseudo-labeled documents, we conduct an ablation experiment. We remove the facility location function and directly select m documents based on the similarity score between external documents and seed documents. The other steps are consistent with FastClass, i.e., it also uses maximum entropy to select the optimal subset. The comparison results are shown in Table 7. We selected four representative tasks: AGnews, DBpedia, Yelp, and SST. The results confirmed our hypothesis: facility location subset selection helped improve model performance in most cases by choosing a more diverse data subset.

Method	AGnews	Dbpedia	SST	Yelp
FastClass _{wo}	84.7	90.3	80.4	83.1
FastClass	85.1	92.4	84.5	87.9

Table 7: Impact of the facility location subset selection on model performance (%). The best performance of each column is in bold. FastClass_{wo}: FastClass without facility location subset selection.

6 Conclusion

We proposed a weakly-supervised text classification method FastClass which selects a document subset returned by dense retrieval models for classifier training. Compared to keyword-driven methods, since our approach no longer needs to expand each class description into a set of class-specific keywords, it is less reliant on initial class descriptions. Extensive experiments show that the proposed method is able to achieve competitive classification performance which does not require high-quality initial class descriptions and often saves significant training time.

Acknowledgement

The authors would like to thank the anonymous referees for their valuable comments. This work is supported by the National Natural Science Foundation of China (No. 61976102 and No. U19A2065), the Science and Technology Development Program of Jilin Province (No.20210508060RQ), and the Fundamental Research Funds for the Central Universities, JLU.

Limitations

The limitations of our method are mainly the following. First, we propose a practically efficient approach without using sophisticated architecture, so the technical approach is relatively simple. Second, the performance of FastClass is competitive on all tasks but only achieves the best results on some tasks.

References

Yuki Arase, Phil Blunsom, Mona Diab, Jesse Dodge, Iryna Gurevych, Percy Liang, Colin Raffel, Andreas Rücklé, Roy Schwartz, Noah A. Smith, Emma Strubell, and Yue Zhang. 2021. [Efficient nlp policy document](#).

- Yoram Baram, Ran El Yaniv, and Kobi Luz. 2004. Online choice of active learning algorithms. *Journal of Machine Learning Research*, 5(Mar):255–291.
- K. Bhatia, K. Dahiya, H. Jain, P. Kar, A. Mittal, Y. Prabhu, and M. Varma. 2016. [The extreme classification repository: Multi-label datasets and code](#).
- John Blitzer, Mark Dredze, Fernando Pereira, and Bollywood Biographies. 2007. Boom-boxes and blenders: domain adaptation for sentiment classification. In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics (ACL'07), Pereira*, volume 447.
- Ming-Wei Chang, Lev-Arie Ratinov, Dan Roth, and Vivek Srikumar. 2008. Importance of semantic representation: Dataless classification. In *Aaai*, volume 2, pages 830–835.
- Zewei Chu, Karl Stratos, and Kevin Gimpel. 2020. Natcat: Weakly supervised text classification with naturally annotated datasets. *arXiv preprint arXiv:2009.14335*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Edwin T Jaynes. 1957. Information theory and statistical mechanics. *Physical review*, 106(4):620.
- Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2017. Billion-scale similarity search with gpus. *arXiv preprint arXiv:1702.08734*.
- Ken Lang. 1995. Newsweeder: Learning to filter net-news. In *Machine Learning Proceedings 1995*, pages 331–339. Elsevier.
- Jens Lehmann, Robert Isele, Max Jakob, Anja Jentzsch, Dimitris Kontokostas, Pablo N Mendes, Sebastian Hellmann, Mohamed Morsey, Patrick Van Kleef, Sören Auer, et al. 2015. Dbpedia—a large-scale, multilingual knowledge base extracted from wikipedia. *Semantic web*, 6(2):167–195.
- Ximing Li, Changchun Li, Jinjin Chi, Jihong Ouyang, and Chenliang Li. 2018. Dataless text classification: A topic modeling approach with document manifold. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, pages 973–982.
- Ximing Li and Bo Yang. 2018. A pseudo label based dataless naive bayes algorithm for text classification with seed words. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1908–1917.
- Han Liu, Xiaotong Zhang, Lu Fan, Xuandi Fu, Qimai Li, Xiao-Ming Wu, and Albert YS Lam. 2019a. Reconstructing capsule networks for zero-shot intent classification. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4799–4809.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019b. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Kyle Lo, Lucy Lu Wang, Mark Neumann, Rodney Kinney, and Daniel Weld. 2020. [S2ORC: The semantic scholar open research corpus](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4969–4983, Online. Association for Computational Linguistics.
- Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Dheeraj Mekala and Jingbo Shang. 2020. Contextualized weak supervision for text classification. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 323–333.
- Yu Meng, Jiaming Shen, Chao Zhang, and Jiawei Han. 2018. Weakly-supervised neural text classification. In *proceedings of the 27th ACM International Conference on information and knowledge management*, pages 983–992.
- Yu Meng, Jiaming Shen, Chao Zhang, and Jiawei Han. 2019. Weakly-supervised hierarchical text classification. In *Proceedings of the AAAI conference on artificial intelligence*, 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 *EMNLP (1)*.
- Jinseok Nam, Eneldo Loza Mencía, and Johannes Fürnkranz. 2016. All-in text: Learning document, label, and word representations jointly. In *Thirtieth AAAI Conference on Artificial Intelligence*.
- George L Nemhauser, Laurence A Wolsey, and Marshall L Fisher. 1978. An analysis of approximations for maximizing submodular set functions—i. *Mathematical programming*, 14(1):265–294.
- Laura Ana Maria Oberländer and Roman Klinger. 2018. An analysis of annotated corpora for emotion classification in text. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2104–2119.
- Pushpankar Kumar Pushp and Muktabh Mayank Srivastava. 2017. Train once, test anywhere: Zero-shot learning for text classification. *arXiv preprint arXiv:1712.05972*.

- Nils Reimers and Iryna Gurevych. 2019. *Sentence-bert: Sentence embeddings using siamese bert-networks*. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Anthony Rios and Ramakanth Kavuluru. 2018. Few-shot and zero-shot multi-label learning for structured label spaces. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*. *Conference on Empirical Methods in Natural Language Processing*, volume 2018, page 3132. NIH Public Access.
- Jacob M Schreiber, Jeffrey A Bilmes, and William Stafford Noble. 2020. apricot: Submodular selection for data summarization in python. *J. Mach. Learn. Res.*, 21:161–1.
- Roy Schwartz, Jesse Dodge, Noah A Smith, and Oren Etzioni. 2020. Green ai. *Communications of the ACM*, 63(12):54–63.
- Jiaming Shen, Wenda Qiu, Yu Meng, Jingbo Shang, Xiang Ren, and Jiawei Han. 2021. Taxoclass: Hierarchical multi-label text classification using only class names. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4239–4249.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 1631–1642.
- Yangqiu Song and Dan Roth. 2014. On dataless hierarchical text classification. In *Twenty-eighth AAAI conference on artificial intelligence*.
- Nandan Thakur, Nils Reimers, Andreas Rücklé, Abhishek Srivastava, and Iryna Gurevych. 2021. Beir: A heterogenous benchmark for zero-shot evaluation of information retrieval models. *arXiv preprint arXiv:2104.08663*.
- Wei Wang, Vincent W Zheng, Han Yu, and Chunyan Miao. 2019. A survey of zero-shot learning: Settings, methods, and applications. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 10(2):1–37.
- Zihan Wang, Dheeraj Mekala, and Jingbo Shang. 2021. X-class: Text classification with extremely weak supervision. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3043–3053.
- Kai Wei, Rishabh Iyer, and Jeff Bilmes. 2015. Submodularity in data subset selection and active learning. In *International Conference on Machine Learning*, pages 1954–1963. PMLR.
- Fangzhao Wu, Ying Qiao, Jiun-Hung Chen, Chuhan Wu, Tao Qi, Jianxun Lian, Danyang Liu, Xing Xie, Jianfeng Gao, Winnie Wu, et al. 2020. Mind: A large-scale dataset for news recommendation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3597–3606.
- Congying Xia, Chenwei Zhang, Xiaohui Yan, Yi Chang, and Philip S Yu. 2018. Zero-shot user intent detection via capsule neural networks. *arXiv preprint arXiv:1809.00385*.
- Zhiqian Ye, Yuxia Geng, Jiaoyan Chen, Jingmin Chen, Xiaoxiao Xu, SuHang Zheng, Feng Wang, Jun Zhang, and Huajun Chen. 2020. Zero-shot text classification via reinforced self-training. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3014–3024.
- Wenpeng Yin, Jamaal Hay, and Dan Roth. 2019. Benchmarking zero-shot text classification: Datasets, evaluation and entailment approach. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3914–3923.
- Rowan Zellers, Ari Holtzman, Hannah Rashkin, Yonatan Bisk, Ali Farhadi, Franziska Roesner, and Yejin Choi. 2019. Defending against neural fake news. In *Advances in Neural Information Processing Systems 32*.
- 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.
- Lu Zhang, Jiandong Ding, Yi Xu, Yingyao Liu, and Shuigeng Zhou. 2021. Weakly-supervised text classification based on keyword graph. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 2803–2813.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. *Advances in neural information processing systems*, 28:649–657.

A Class Descriptions in Evaluation Datasets

We list the class descriptions of the datasets we used for evaluation as follows. These texts are used as to compute SBERT vector representations.

Yahoo: Society and Culture; Science and Mathematics; Health, Education and Reference; Computers and Internet; Sports; Business and Finance; Entertainment and Music; Family and Relationships; Politics and Government.

AGnews: politics; sports; business; technology.

20Newsgroup: atheist atheism; computer graphics; computer OS microsoft windows miscellaneous; computer system IBM PC hardware; computer system Mac hardware; computer windows; miscellaneous for sale; recreational automobile; recreational motorcycles; recreational sport baseball; recreational sport hockey; science cryptography; science electronics; science medicine; science space; society religion christian; talk politics guns; talk politics middle East; talk politics miscellaneous; talk religion miscellaneous.

DBPedia: Company; Educational Institution; Artist; Athlete; Office Holder; Mean Of Transportation; Building; Natural Place; Village; Animal; Plant; Album; Film; Written Work.

Yelp: positive; negative.

Amazon: positive; negative.

SST-B: positive; negative.

Emotion: anger; sadness; surprise; love; fear; disgust; guilt; shame; joy; no emotion.

Situation: utilities energy or sanitation; water supply; search/rescue; medical assistance; infrastructure; shelter; evacuation; regime change; food supply; crime violence; terrorism; out-of-domain.

B BERT-based Classifier Performance

We have reported the results based on RoBERTa as our main result. Here we show the classification performance based on BERT classifier in Table 8. In most cases, we found that the performance of RoBERTa model is better than BERT. This may be because compared with BERT’s use of Wikipedia and books the training data of RoBERTa comes from web text which is more diverse.

C Relation Between Entropy and Accuracy

In order to verify the relationship between entropy and classification accuracy, we compared the trends

of entropy and predicate accuracy under different parameter settings. Figure 4 shows the relation between the entropy and accuracy in evaluation datasets. From Figure 4 we can see that with different parameters, the trends of entropy and accuracy are often (but not perfectly) correlated. It shows that the empirical classification entropy on unlabeled data is an effective unsupervised metric to guide the selection of pseudo-labeled subset.

Method	Yahoo	AGnews	20News	DBPedia	Yelp	Emotion	Amazon	SST	Situation
FastClass _{Bert}	60.8	83.9	53.9	92.8	86.7	26.8	84.1	81.8	38.0
FastClass _{RoBerta}	61.9	85.1	55.4	92.4	87.9	28.0	89.1	84.5	39.2

Table 8: Weakly-supervised text classification performance in nine datasets (%) based on BERT and RoBERTa classifier.

