

CEIL: A General Classification-Enhanced Iterative Learning Framework for Text Clustering

Mingjun Zhao*
University of Alberta
Edmonton, Canada
DiDi Global
Beijing, China
zhao2@ualberta.ca

Mengzhen Wang*
DiDi Global
Beijing, China
wangmengzhen_i@didiglobal.com

Yinglong Ma
North China Electric Power
University
Beijing, China
yinglongma@ncepu.edu.cn

Di Niu
University of Alberta
Edmonton, Canada
dniu@ualberta.ca

Haijiang Wu✉
DiDi Global
Beijing, China
wuhaijiang@didiglobal.com

ABSTRACT

Text clustering, as one of the most fundamental challenges in unsupervised learning, aims at grouping semantically similar text segments without relying on human annotations. With the rapid development of deep learning, deep clustering has achieved significant advantages over traditional clustering methods. Despite the effectiveness, most existing deep text clustering methods rely heavily on representations pre-trained in general domains, which may not be the most suitable solution for clustering in specific target domains. To address this issue, we propose *CEIL*, a novel Classification-Enhanced Iterative Learning framework for short text clustering, which aims at generally promoting the clustering performance by introducing a classification objective to iteratively improve feature representations. In each iteration, we first adopt a language model to retrieve the initial text representations, from which the clustering results are collected using our proposed Category Disentangled Contrastive Clustering (CDCC) algorithm. After strict data filtering and aggregation processes, samples with clean category labels are retrieved, which serve as supervision information to update the language model with the classification objective via a prompt learning approach. Finally, the updated language model with improved representation ability is used to enhance clustering in the next iteration. Extensive experiments demonstrate that the *CEIL* framework significantly improves the clustering performance over iterations, and is generally effective on various clustering algorithms. Moreover, by incorporating *CEIL* on CDCC, we achieve the state-of-the-art clustering performance on a wide range of short text clustering benchmarks outperforming other strong baseline methods.

CCS CONCEPTS

• Computing methodologies → Information extraction; Natural language processing; Artificial intelligence.

KEYWORDS

Text Clustering, Classification-enhanced Clustering, Iterative Framework

ACM Reference Format:

Mingjun Zhao, Mengzhen Wang, Yinglong Ma, Di Niu, and Haijiang Wu✉. 2023. CEIL: A General Classification-Enhanced Iterative Learning Framework for Text Clustering. In *Proceedings of the ACM Web Conference 2023 (WWW '23)*, May 1–5, 2023, Austin, TX, USA. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3543507.3583457>

1 INTRODUCTION

Text clustering, as an unsupervised approach to organizing semantics and analyzing data distributions, plays a vital role in many data-driven natural language processing tasks, such as news recommendation [12, 32], topic extraction [31], document summarization [21], etc.

Traditional text clustering approaches transform text into high-dimensional representations based on numerical statistics, such as bag-of-words (BOW) and term frequency-inverse document frequency (TF-IDF), after which a clustering algorithm is applied to partition text segments into homogeneous groups with respect to a given distance function. However, in the application of short text clustering, text representations created with traditional measures tend to suffer from sparsity issues, due to the limited length and information within each text piece.

Motivated by the huge success of deep learning in recent years, deep clustering approaches have been increasingly popular. With the utilization of low-dimensional dense representations extracted from neural networks, deep clustering methods better capture the semantic meaning and alleviate the sparsity issue. Due to unavailability of supervision in target datasets, unsupervised tasks such as the reconstruction task are often utilized for representation learning with an auto-encoder network [7, 24, 26].

Many studies also leverage the large-scale general domain data for better feature representations. Word2Vec [13] and GloVe [15] learn an embedding vector for each word and are popular choices for text clustering [4, 25]. Recently, pre-trained language models (PLMs) such as BERT [5] and GPT [17] have achieved great success in improving the representation ability of models, which leads to better clustering results [25]. Despite the success, PLMs tend to produce homogeneous representations for all text samples in different clusters, which limits the clustering performance. Moreover, directly relying on PLM representations may not be optimal due to

*Both authors contributed equally to this research.

the gap between the general-purpose corpus on which PLMs are pretrained and the target clustering dataset.

In this paper, we propose *CEIL*, a novel Classification-Enhanced Iterative Learning framework generally applicable to various short text clustering algorithms, in which a classification objective is introduced to improve text representations. The clustering task and the introduced classification task iteratively promote each other, leading to significant performance improvements.

Specifically, in each iteration, we first extract feature representations with a classification-enhanced language model derived from last iteration. We use a prompt learning approach to extract text representations, as it can better capture the semantics with limited supervision data compared to the fine-tuning approach. Then, we propose Category Disentangled Contrastive Clustering (CDCC) algorithm, where a novel category loss is designed to encourage the separation of different clusters, producing better clustering results with more explicit boundaries. The results are further processed with strict data filtering and aggregation algorithms to alleviate the influence of improperly clustered samples, and are transformed into pseudo classification labels using an automatic verbalizer. Finally, by learning a classification objective on the clustering dataset using pseudo labels, the language model is updated and acquires better representation ability.

Extensive experiments demonstrate that our *CEIL* framework is generally effective when used with various clustering algorithms, and improves the performance over iterations by exploiting the classification objective to obtain better text representations. By incorporating *CEIL* on our proposed CDCC algorithm, we achieve a best average accuracy of 83.2% over 9 short text clustering benchmarks, significantly outperforming other competitive baseline methods by a substantial margin.

2 RELATED WORK

The process of short text clustering mainly involves the derivation of text representations and the clustering framework. In this section, we first provide a brief overview of text representation methods, and then introduce different clustering frameworks.

Text Representations. Compared with the general text clustering problem, the task of short text clustering is more challenging due to the limited number of words in each text piece, where traditional representation methods such as BoW and TF-IDF often result in sparse vectors with poor expressive ability. Several studies have addressed this issue by enriching the short text segments with external resources. Banerjee et al. [2], Hu et al. [9] propose to map the short text segments to Wikipedia articles and leverage the background knowledge for better clustering performance. In a similar fashion, ontology hierarchical relations and lexical chains are utilized to enhance text representations for clustering [22].

Low-dimensional representations, such as word embeddings, learned with neural networks have been widely explored and shown great potential to alleviate the sparsity issue. Different aggregation methods have been explored to produce the sentence representations from word embeddings, including mean-pooling, max-pooling, and weighted summation based on frequency information [4]. Although these methods are simple, the word order is lost, which may result in inaccurate semantics. Recurrent neural networks are

exploited to overcome this problem by sequentially modeling the input words and deriving an integrated sentence representation [10].

The recent advances of pre-trained language models have achieved great success in building text representations by modeling contextual information [5, 17], and shown considerable improvements on short text clustering [11, 29].

Clustering Framework. Most of the studies in short text clustering adopt a simple framework which focuses on the representation learning, and apply clustering algorithms on top of the well-learned representations. In DEC [24] and Self-Train [7], auto-encoders are employed to solve the reconstruction tasks with the aim of enhancing text representations. STC² [25] incorporates dimensional reduction methods to create auxiliary labels and adopts a convolutional neural network to learn text representations by reconstructing these auxiliary targets. In SCCL [29], authors leverage contrastive learning objectives to promote better separated representations. Despite the promising performance, there's still room for improvement due to the discrepancy between the representation learning objective and the clustering objective.

Another design idea of clustering framework is to iteratively adjust existing clustering results. Rakib et al. [18] propose to train a classification model with cluster labels to correct an initial clustering result by reclassifying the outliers. However, this method uses fixed representations for clustering, thus may not generalize well.

The proposed *CEIL* framework aims to combine the advantages of both designs, by iteratively improving the clustering performance with enhanced text representations, and leveraging the more accurate clusters to promote the representations with the introduced classification objective. Benefiting from this delicate design, the *CEIL* framework is capable of producing more accurate clusters and steadily improving the clustering performance, compared to previous approaches.

3 METHODOLOGY

In this section, we first provide a general overview of the proposed *CEIL* framework, and introduce our design for the clustering and the classification in detail.

3.1 Overall Architecture

Figure 1 gives an overview of our proposed framework, in which a clustering objective as well as a classification objective are designed to promote each other iteratively. In each iteration, the clustering process provides supervisory category labels for classification, and the classification task updates the language model to produce better text representations for clustering.

Specifically, in clustering, we adopt a pre-trained language model at the first iteration, and use the language model trained on the classification task in the following iterations, to extract the text representations. The prompt-tuning approach is applied during representation extraction, as it is able to produce more semantically meaningful representations especially under the circumstances when the downstream data is limited where little fine-tuning can be performed. Then, with the extracted representations, we can apply any applicable clustering algorithm, e.g., simple K-means or our proposed CDCC, to derive the text clusters. The clustering

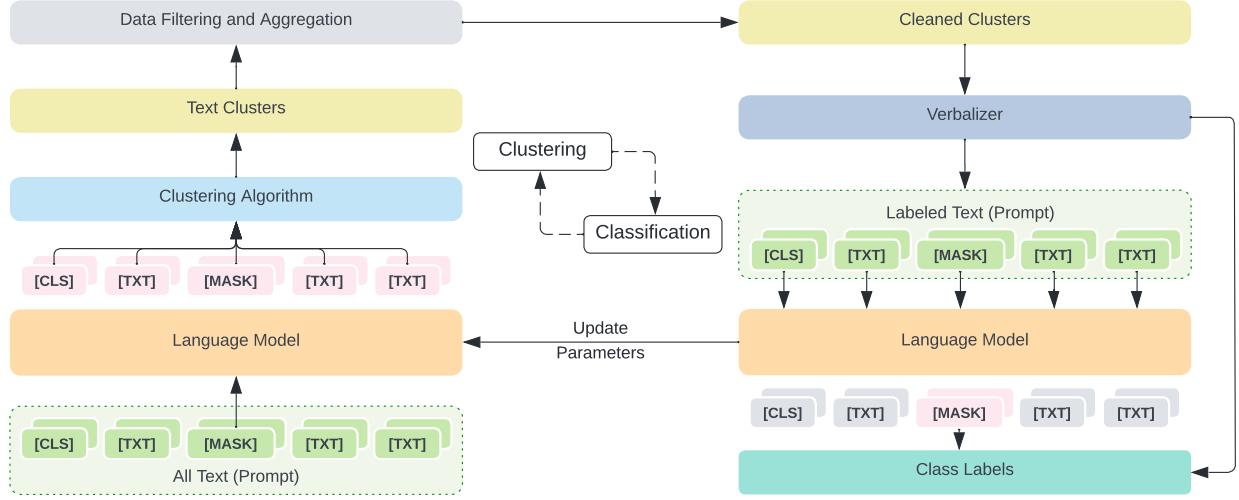


Figure 1: An overview of the proposed Classification-Enhanced Iterative Learning (CEIL) framework for short text clustering. In each iteration of CEIL, the clustering task and the classification task are alternatively optimized. The clustering exploits the classification-enhanced language model for better text representations and outputs more precise text clusters. After data filtering and aggregation, the improved clusters are processed with a verbalizer to extract more accurate classification labels, which again enhances the representation ability of the language model by optimizing the classification objective.

results are cleaned with strict filtering and aggregation to reduce the influence of inaccurate clustering. Then, the cleaned results are processed with a verbalizer to derive accurate category labels which serve as supervisory signals for the classification task. By training to classify the text segments into different categories, we effectively enhance the representation ability of the language model, which further benefits the clustering in the next iteration.

3.2 Category Disentangled Contrastive Clustering

Contrastive Learning (CL) [3, 8, 23] aims at learning low-dimensional representations of data where similar samples are pulled together, while dissimilar samples are pushed apart. In SCCL [29], a contrastive loss is incorporated along with the clustering loss to encourage better separation of representations. Despite its effectiveness, the contrastive loss encourages separation of any pair of two samples, disregarding their semantic similarities. As a potential result, we observe that many samples locate at the halfway of two clusters, leading to ambiguous cluster boundaries.

In order to resolve this issue, we propose Category Disentangled Contrastive Clustering (CDCC) which improves the existing CL-based clustering algorithm by combining the existing contrastive loss and clustering loss with a novel category objective to encourage better separation of different clusters and more explicit cluster boundaries.

Representation Extraction. Either traditional clustering or deep clustering algorithms require the representations of data samples as

input. The quality of initial data representations are crucial to clustering and have great impact on the final performance.

Pre-trained language models show great potential on constructing contextual text representations. However, without fine-tuning on target data, the expressive ability of the produced representations is limited. In order to obtain better initial text representations, we adopt the prompt approach for extracting representations by appending a carefully designed template with a [MASK] token, which better summarizes the semantics of the text. Specifically, in the first iteration of CEIL, as the language model has not been updated with prompt tuning, we choose the mean vector of all tokens in the prompted input text as the representation, while in the following iterations, the vector of [MASK] produced from the PLM is used as the representation.

Contrastive Loss. Specifically, given a mini-batch of M input text segments $\mathcal{B} = \{x_k\}_{k=1}^M$, an augmented batch $\mathcal{B}^a = \{\tilde{x}_k\}_{k=1}^{2M}$ is constructed by applying two random augmentations on each sample x_k to derived augmented samples \tilde{x}_{2k-1} and \tilde{x}_{2k} as a positive pair in contrastive learning. The other samples in \mathcal{B}^a are treated as negative instances regarding this positive pair. Let $\tilde{h}_i = \phi(\tilde{x}_i)$ be the representation of sample \tilde{x}_i extracted with language model ϕ . The loss function for a positive pair of samples (i, j) is defined as

$$\ell(i, j) = -\log \frac{\exp(\cos(\tilde{h}_i, \tilde{h}_j)/\tau)}{\sum_{k=1}^{2M} \mathbb{1}_{k \neq i} \cdot \exp(\cos(\tilde{h}_i, \tilde{h}_k)/\tau)}, \quad (1)$$

where $\cos(\cdot)$ refers to the cosine similarity function. The contrastive loss is derived by calculating the average loss on all samples in the




Figure 2: The data filtering process. (a) The sample distribution in a cluster. (b) The generated graph. (c) The search procedure. (d) The filtered cluster.

augmented batch \mathcal{B}^a :

$$\mathcal{L}_{\text{contrast}} = \frac{1}{2M} \sum_{i=1}^M [\ell(2k-1, 2k) + \ell(2k, 2k-1)]. \quad (2)$$

Clustering Loss. The clustering loss is first proposed in DEC [24], and adopted in short text clustering [7, 29]. It aims to improve the clustering by iterating between computing an auxiliary target distribution and minimizing the Kullback-Leibler (KL) divergence to it.

Given a text sample x_i in mini-batch \mathcal{B} , its representation h_i , and the learnable centroid vector μ_j of cluster j , we calculate the probability q_{ij} of assigning x_i to cluster j with the Student's t -distribution:

$$q_{ij} = \frac{(1 + \|h_i - \mu_j\|_2^2 / \alpha)^{-\frac{\alpha+1}{2}}}{\sum_{j'=1}^K (1 + \|h_i - \mu_{j'}\|_2^2 / \alpha)^{-\frac{\alpha+1}{2}}}, \quad (3)$$

where K is the number of clusters and α denotes the degree of freedom. Following Van der Maaten and Hinton [20], we set α to 1.

The auxiliary distribution p_{ij} is computed as:

$$p_{ij} = \frac{q_{ij}^2 / f_i}{\sum_{j'} q_{ij'}^2 / f_{j'}}, \quad (4)$$

where $f_j = \sum_i q_{ij}$ is the soft cluster frequencies and is approximated within the M samples in mini-batch \mathcal{B} .

The clustering loss is defined to jointly optimize the cluster centroids $\{\mu_j\}_{j=1}^K$ and the language model ϕ by minimizing the KL-divergence between the soft assignments q_i and the auxiliary distribution p_i as:

$$\mathcal{L}_{\text{cluster}} = \sum_{i=1}^M \text{KL}[p_i \| q_i] = \sum_{i=1}^M \sum_{j=1}^K p_{ij} \log \frac{p_{ij}}{q_{ij}}. \quad (5)$$

Category Loss. One critical problem of clustering is that the initial data representations have significant overlap across categories. Although the fore-mentioned contrastive learning objective alleviates this issue by separating individual samples in the representation space, we still observe the presence of many sample points located at the shared boundary of two clusters, potentially leading to incorrect clustering and ambiguous cluster boundaries. To resolve this issue, we propose a novel category disentangling objective which aims at improving clustering by introducing more room between clusters and obtaining more explicit cluster boundaries.

Based on the probability q_{ij} calculated with Equation 3, we first make hard assignments for each sample x_i to cluster c_k , where $k = \text{argmax}_j(q_{ij})$. Then, for each cluster c_k , we obtain its initial representation vector $h_{c_k} = \text{avg}(\{x_i | x_i \in c_k\})$ by taking the average of the representations of samples in it.

The initial representation can be inaccurate especially during the first few iterations when the clustering results are not very stable. Therefore, we apply a simple threshold θ to obtain a filtered cluster $c'_k = \{x_i | \cos(h_i, h_{c_k}) \geq \theta, x_i \in c_k\}$. Similarly, for each filtered cluster c'_k , we derive a final representation vector $h_{c'_k} = \text{avg}(\{x_i | x_i \in c'_k\})$. With the aim of better separating the clusters, the category loss is proposed to push apart the cluster feature representations by maximizing the cosine distance between clusters, which is defined as:

$$\mathcal{L}_{\text{category}} = \exp\left(\frac{\sum_{i=1}^{K-1} \sum_{j=i+1}^K \cos(h_{c'_i}, h_{c'_j})}{(K-1)(K-2)/2}\right). \quad (6)$$

Overall Objective. In summary, the overall objective of CDCC is:

$$\mathcal{L} = \mathcal{L}_{\text{contrast}} + \lambda \mathcal{L}_{\text{cluster}} + \mathcal{L}_{\text{category}}, \quad (7)$$

where λ is a loss coefficient to balance the objectives.

3.3 Data Filtering and Aggregation

After clustering, each sample x_i is assigned a category label c_k , which can later be utilized as the supervisory label for the classification objective. However, the results produced by clustering algorithm can hardly be perfect, where samples from multiple ground truth categories can be allocated into one cluster, and samples from the same category may be split into multiple clusters. Directly incorporating the unprocessed clustering results as training data for classification may introduce much noise and jeopardize the performance of the classifier. Therefore, we propose to incorporate a data filtering as well as a data aggregation module to alleviate this problem.

Data Filtering. We propose a data filtering algorithm for cleaning clustering results by mining the adjacency relationships between sample points in the same cluster and keeping a collection of samples with abounding inter-relations as the filtered cluster.

The data filtering process is illustrated in Figure 2, where the sample distribution of a cluster is shown in (a). In (b), we create a graph where the nodes represent the samples and the edges between nodes represent the adjacency relationship, which are created when

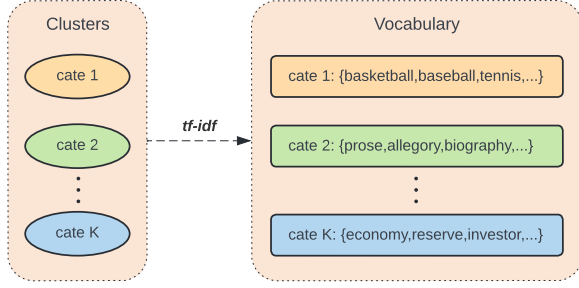


Figure 3: Our automatic verbalizer maps the cluster categories to correlated tokens in the vocabulary based on token frequencies.

the cosine similarity of two nodes is greater than a threshold β . The search procedure is demonstrated in (c). Starting from the center node colored in red with the most connections, we retrieve its adjacent neighbors colored in green and itself through the red arrows. The same process is repeated on all green nodes to retrieve the second-order neighbors colored in yellow through the green arrows. We take all the retrieved nodes as candidates, gather the statistics of their input degrees, i.e., the number of arrows pointing to a node, and select the samples whose input degree is greater than the mode value as the filtered cluster as shown in (d).

This data filtering process guarantees only tightly clustered samples remain, which efficiently alleviate the noise that could be potentially introduced to the classification task.

Data Aggregation. This process aims to merge clusters close enough to each other which are very likely to come from the same ground truth category. We adopt a simple aggregation strategy by measuring the similarity between two clusters c'_i and c'_j , and merging these two clusters if their similarity is greater than a threshold δ . The merge will eliminate the original two clusters and produce a new cluster $c''_{ij} = c'_i \cup c'_j$.

Empirically, the data aggregation proves to be effective in improving the clustering performance especially on imbalanced datasets with plenty of categories.

3.4 Classifier Verbalizer

In prompt learning, the verbalizer is an important module that bridges a projection between the vocabulary and class labels, and has critical impacts on the classification performance. Most existing works use manual verbalizers, where specific tokens are selected by human for each class, which is not only time-consuming and requiring expert knowledge, but also not applicable to our clustering circumstances, where no supervision is available.

As shown in Figure 3, we propose an automatic verbalizer to map meaningless category labels in clustering to concrete representative tokens in the vocabulary. Specifically, for each cleaned cluster c''_i , we gather the token frequencies and extract a collection of n keywords $\mathcal{V}_{c''_i} = \{t_k\}^n$ based on the TF-IDF value of each token in the

Table 1: Dataset statistics. $|\mathcal{V}|$: the vocabulary size; N: number of short text documents; L: average number of words in each document; C: number of clusters; P: the ratio of the size of the largest cluster to that of the smallest cluster.

Dataset	$ \mathcal{V} $	N	L	C	P
AgNews	21K	8000	23	4	1
StackOverflow	15K	20000	8	20	1
Biomedical	19K	20000	13	20	1
SearchSnippets	31K	12340	18	8	7
GoogleNews-TS	20K	11109	28	152	143
GoogleNews-S	18K	11109	22	152	143
GoogleNews-T	8K	11109	6	152	143
Tweet	5K	2472	8	89	249
MX	17K	20000	26	210	15

vocabulary \mathcal{V} . The extracted keywords $\mathcal{V}_{c''_i}$ are then utilized as the supervisory labels for the classification task.

3.5 Classification

Inspired by [27], we use the classification objective to update the language model ϕ for better representation learning, and further improve the clustering performance.

Consistent with the update of language model in clustering, we also adopt prompt learning for the optimization of the classification task. Given a prompted text sample x_j in cluster c''_i and the corresponding classification labels $\mathcal{V}_{c''_i}$, the objective is to maximize the probability of the label y_j calculated with the language model ϕ by averaging the probabilities of tokens in $\mathcal{V}_{c''_i}$ filled in the [MASK] token:

$$\begin{aligned} & \max_{\phi} P_{\phi}(y_j | x_j) \\ & = \max_{\phi} \left[\frac{1}{|\mathcal{V}_{c''_i}|} \sum_{t \in \mathcal{V}_{c''_i}} (P_{\phi}([MASK] = t | x_j)) \right], \end{aligned} \quad (8)$$

where the update of the language model ϕ is achieved by gradient descent.

By utilizing the strictly cleaned clustering results as the supervisory signals, the representation ability of the language model is further improved by learning on the classification objective to capture the semantics of the text. The updated language model is then exploited in the next iteration for better text representations, leading to more accurate clusters and classification labels of higher quality, which again improves the language model.

With the iterative update on the clustering and classification objectives, and the fore-mentioned delicately designed modules, our CEIL framework learns better text representations and significantly improves clustering performance.

4 EXPERIMENTS

In this section, we will first describe the datasets used in our evaluation and provide the implementation details along with the performance results.

Table 2: The hyperparameters and prompt templates adopted for different tasks in our experiment.

Dataset	lr	$bs_{cluster}$	$bs_{classifier}$	θ	β	δ	n	Prompt Template
AgNews	1e-5	100	64	0.7	0.93	1	15	[X] This topic is about [MASK] .
StackOverflow	1e-5	400	300	0.7	0.6	1	24	[X] [MASK] .
Biomedical	1e-5	100	64	0.7	0.6	1	20	[MASK] [X] .
SearchSnippets	1e-5	100	64	0.7	0.6	1	25	[X] is [MASK] .
GoogleNews-TS	5e-6	100	64	0.7	0.8	0.98	9	[MASK] [X] .
GoogleNews-S	1e-5	100	64	0.7	0.76	0.99	39	[X] This topic is about [MASK] .
GoogleNews-T	5e-6	400	300	0.7	0.78	0.98	35	Category : [MASK] [X]
Tweet	1e-5	100	64	0.7	0.8	0.95	10	[MASK] [X] .
MX	1e-5	100	64	0.7	0.6	0.95	10	This food is made up of [MASK] .

Datasets. Our experiments are based on eight public short text clustering datasets and an internal dataset MX. Table 1 summarizes the statistics of the evaluation datasets.

- **AgNews** [30] is a dataset of news articles. The news titles of 4 topics selected by HAC-SD [18] are used as the short text segments for clustering.
- **StackOverflow** ¹ contains 20,000 question titles associated with 20 categories from StackOverflow websites, selected by STCC [25].
- **Biomedical** ² is a subset of the PubMed data distributed by BioASQ. It contains randomly selected 20,000 page titles from 20 groups by STCC [25].
- **SearchSnippets** [16] is a dataset extracted from web search snippets, consists of 12,340 snippets with 8 groups.
- **GoogleNews** [28] contains news titles and snippets from 11,109 articles associated with 152 events. Following HAC-SD [18], GoogleNews-TS, GoogleNews-T, and GoogleNews-S are three subsets containing titles and snippets, titles only, and snippets only, respectively.
- **MX** is an internal dataset on dish recipes in Spanish, consists of more than 3M text in 150 course-grained categories, which are split into fine-grained clusters. Table 1 gives the average statistics on different categories.

Evaluation Metrics. Following the common practice, we evaluate the performance of different short text clustering methods via top-1 accuracy (ACC) and Normalized Mutual Information (NMI).

Baselines. We compare our proposed method with traditional text clustering methods and a number of state-of-the-art methods, including BoW, TF-IDF, GMM, Self-Train [7], HAC-SD [18], and STCC [29].

- **BoW & TF-IDF** features are extracted with a feature dimension of 1,500, and are incorporated with K-means clustering.
- **GMM** is a traditional clustering method using Gaussian Mixture Models.
- **STCC** [25] incorporates Word2Vec [13] word embeddings and a convolutional neural network to construct the text representation, and applies K-means for clustering.

- **Self-Train** [7] enhances word embeddings using SIF [1], and adopts an auto-encoder for representations.
- **HAC_SD** [18] uses hierarchical agglomerative clustering clustering on a sparse pairwise similarity matrix. **HAC_SD_IC** further updates the clustering results by iteratively reclassifying the outliers.

Implementation Details. We implement our *CEIL* framework based on PyTorch [14] and OpenPrompt [6], and incorporate *distilbert-base-nli-stsb-mean-tokens* from Sentence Transformer [19] as the initial checkpoint of the language model. In the application of *CEIL* framework, we adopt 5 iterations of update when applied on SCCL and CDCC, and 8 iterations on GMM.

In Table 2, we list the adopted hyperparameters and prompt templates for each dataset. We choose the learning rate between 1e-5 and 5e-6, the batch size of clustering between 100 and 400, and the batch size of classifier between 64 and 300. The loss coefficient for the clustering loss λ is set to 1. The threshold θ used in the category loss is set to 0.7. For the choice of filtering threshold β , and aggregation threshold δ , we use a grid search with an interval of 0.01 from 0.6 to 1.0 and 0.95 to 1.0 respectively. And the number of keyword n in verbalizer is also chosen by grid search with an interval of 1. For each dataset, we adopt the best prompt template selected from the classic manual designed templates from OpenPrompt [6] by a simple K-means experiment.

4.1 Comparison with State-of-the-art Baselines

We conduct comprehensive experiments on the evaluation datasets to demonstrate the effectiveness of our proposed CDCC clustering algorithm and the *CEIL* framework. We have also proved the general applicability of the *CEIL* framework by applying it on existing text clustering algorithms including GMM, SCCL and CDCC. For evaluation, we conduct five repetitive runs and report the mean performance in Table 3a, and include the standard deviation results in 3b.

From the results we can see that among the clustering algorithms without iterative updates, our CDCC achieves the best average performance of 78.8% ACC and 76.8% NMI, outperforming the competitive baselines including SCCL [29] and HAC_SD [18]. Note that HAC_SD_IC also adopts iterative updates by iteratively reclassifying the outliers. Out of 9 evaluation tasks, our CDCC beats SCCL in 8 tasks, demonstrating that our proposed category loss

¹<https://www.kaggle.com/c/predict-closed-questions-on-stack-overflow/>

²<http://participants-area.bioasq.org>

Table 3: Comparison results of different short text clustering methods on nine evaluation datasets. The results of applying our proposed *CEIL* framework on GMM, SCCL, and CDCC as GMM_CEIL, SCCL_CEIL, and CDCC_CEIL are also included.**(a) Accuracy and NMI results.**

Method	Average		AgNews		StackOverflow		Biomedical		SearchSnippets		Tweet		GoogleNews-TS		GoogleNews-S		GoogleNews-T		MX	
	ACC	NMI	ACC	NMI	ACC	NMI	ACC	NMI	ACC	NMI	ACC	NMI	ACC	NMI	ACC	NMI	ACC	NMI	ACC	NMI
BOW	-	-	27.6	2.6	18.5	14.0	14.3	9.2	24.3	9.3	49.7	73.6	57.5	81.9	49.0	73.5	49.8	73.2	-	-
TF-IDF	-	-	34.5	11.9	58.4	58.7	28.3	23.2	31.5	19.2	57.0	80.7	68.0	88.9	61.9	83.0	58.9	79.3	-	-
STCC	-	-	-	-	51.1	49.0	43.6	38.1	77.0	63.2	-	-	-	-	-	-	-	-	-	-
Self-Train	-	-	-	-	59.8	54.8	54.8	47.1	77.1	56.7	-	-	-	-	-	-	-	-	-	-
HAC_SD	76.4	69.3	81.8	54.6	64.8	59.5	40.1	33.5	82.7	63.8	89.6	85.2	85.8	88.0	80.6	83.5	81.8	84.2	80.7	71.6
HAC_SD_IC	82.2	75.2	84.5	59.1	78.7	73.4	47.8	41.3	87.7	71.9	91.5	86.9	92.3	93.2	89.0	90.0	87.2	87.9	81.5	72.8
GMM	58.7	60.2	68.0	33.4	57.3	49.2	38.9	32.9	56.3	32.5	52.8	78.1	64.4	86.7	57.5	80.7	54.9	78.1	78.1	69.8
GMM_CEIL	76.2	74.6	86.6	64.7	75.9	74.7	44.0	43.9	77.9	58.9	87.8	92.3	84.3	92.8	77.9	88.1	70.9	84.8	80.6	71.2
SCCL	78.1	76.8	88.2	68.2	75.5	74.5	46.2	41.5	85.2	71.1	78.2	89.2	89.8	94.9	83.1	90.4	75.8	88.3	80.6	73.2
SCCL_CEIL	81.0	77.9	88.9	69.2	77.8	76.1	47.3	39.9	88.3	73.9	87.3	92.6	91.3	95.1	85.5	91.6	79.6	89.1	83.2	74.0
CDCC	78.8	76.8	88.3	68.2	76.9	75.3	42.8	38.1	85.3	71.2	82.6	90.2	90.2	94.9	84.5	90.6	76.2	88.5	82.8	74.3
CDCC_CEIL	83.2	79.2	89.1	69.7	78.8	77.3	49.2	41.0	89.3	75.2	91.2	94.1	92.5	95.8	87.2	92.4	84.3	91.3	87.3	76.1


(b) Standard deviations results.

Method	AgNews		StackOverflow		Biomedical		SearchSnippets		Tweet		GoogleNews-TS		GoogleNews-S		GoogleNews-T		MX	
	ACC_std	NMI_std	ACC_std	NMI_std	ACC_std	NMI_std	ACC_std	NMI_std	ACC_std	NMI_std	ACC_std	NMI_std	ACC_std	NMI_std	ACC_std	NMI_std	ACC_std	NMI_std
GMM	0.66	0.58	1.30	0.61	0.65	1.31	1.6	0.72	1.30	0.66	0.11	0.22	0.10	0.12	0.69	0.50	0.40	0.13
GMM_CEIL	0.52	0.58	0.92	0.51	0.21	0.31	0.66	0.78	0.98	0.78	0.17	0.09	0.61	0.55	0.71	0.23	0.36	0.28
CDCC	0.23	0.54	1.70	2.05	0.17	2.4	1.30	1.40	0.61	0.39	0.48	0.26	0.33	0.32	0.70	0.32	0.21	0.10
CDCC_CEIL	0.17	0.33	0.70	0.33	1.20	0.57	0.60	1.03	1.96	0.78	0.34	0.08	1.27	0.18	0.83	0.26	0.21	0.13

can effectively obtain better separation of clusters and generally improve the clustering performance under different circumstances.

The *CEIL* framework also shows promising effects on boosting clustering performance. By comparing the *CEIL* enhanced methods GMM_CEIL, SCCL_CEIL and CDCC_CEIL with their original version, we can see that applying *CEIL* framework stably improves the clustering performance on all evaluation tasks, and brings significant performance gains of 17.5%, 2.9% and 4.4% average accuracy improvements on the traditional GMM method, and more advanced deep clustering methods, SCCL and CDCC, respectively. Moreover, among all clustering methods, our CDCC_CEIL achieves the top accuracy on 5 out of 9 tasks and the best NMI on 8 out of 9 tasks, while obtaining the best average performance of 83.2% ACC and 79.2% NMI. This result provides strong evidence that the proposed *CEIL* framework can produce better text clusters by iteratively promoting the text representations with the classification objective, and is generally effective to various text clustering algorithms.

We have also conducted thorough analysis to examine the reason why Self-Train [7] has a clear advantage over other methods on the Biomedical dataset. The Biomedical dataset contains professional domain-specific knowledge which are hardly covered in general pre-training corpus, while Self-Train [7] incorporates an additional large-scale in-domain corpus to pre-train the word embeddings, leading to better clustering performance. From this observation, we can infer a potential improvement to our method of conducting continual pre-training on in-domain datasets to capture domain-specific knowledge before clustering.

**(a) Inter-cluster Distances****(b) Top-1 Accuracy****Figure 4: The GoogleNews-S results of CDCC and SCCL in terms of inter-cluster distances and top-1 accuracy during the learning of clustering.**

4.2 Ablation Studies



We have conducted comprehensive ablation analysis to evaluate the impact of different modules in our method, including the category loss in the proposed CDCC algorithm, the data filtering and aggregation module, the prompt learning approach for representation construction and classification learning, and the iterative updating of the *CEIL* framework.

Category Loss. While the main difference between CDCC and SCCL is the additional category loss incorporated in CDCC, we have conducted experiments on GoogleNews-S to evaluate its effectiveness by comparing the inter-cluster distances and accuracy of CDCC and SCCL over iterations.

Figure 4 shows the comparison results, where the CDCC algorithm produces more separated clusters and constantly outperforms SCCL with higher accuracy by incorporating the proposed category

Table 4: Comparison results of the prompt learning approach (“_p”) and the fine-tuning approach (“_f”) for feature extraction and classification learning.


Method	Average		AgNews		StackOverflow		Biomedical		SearchSnippets		Tweet		GoogleNews-TS		GoogleNews-S		GoogleNews-T		MX	
	ACC	NMI	ACC	NMI	ACC	NMI	ACC	NMI	ACC	NMI	ACC	NMI	ACC	NMI	ACC	NMI	ACC	NMI	ACC	NMI
CDCC_f	78.8	76.8	88.3	68.2	76.9	75.3	42.8	38.1	85.3	71.2	82.6	90.2	90.2	94.9	84.5	90.6	76.2	88.5	82.8	74.3
CDCC_p	79.9	76.7	88.3	68.2	77.0	75.3	43.4	34.2	88.0	72.3	86.8	91.9	90.8	95.0	84.8	90.8	77.1	88.7	83.2	74.5
CDCC_CEIL_f	81.2	77.6	88.5	68.3	77.6	76.5	47.5	38.1	88.2	73.1	88.2	92.1	91.2	95.5	86.8	91.0	78.5	88.3	84.1	75.1
CDCC_CEIL_p	83.2	79.2	89.1	69.7	78.8	77.3	49.2	41.0	89.3	75.2	91.2	94.1	92.5	95.8	87.2	92.4	84.3	91.3	87.3	76.1

**(a) Results w/o data filtering****(b) Results w/o data aggregation****Figure 5: The ablative analysis of the data filtering and aggregation modules in CEIL evaluated using GMM and CDCC algorithms. “_rf” and “_ra” denote the filtering and aggregation module is removed from the clustering framework respectively.**

loss. This observation validates our hypothesis that the category loss achieves better separation between clusters, leading to more accurate clustering.

Data Filtering and Aggregation. Experiments were conducted to assess the impact of the data filtering and aggregation module by separately removing them from the framework and evaluating the clustering performance.

In Figure 5a and 5b, we demonstrate the comparison results of the complete CEIL framework on GMM and CDCC with the frameworks where the data filtering or aggregation module is removed, denoted with the “_rf” and “_ra” suffixes respectively. Results show

**(a) GMM_CEIL results****(b) CDCC_CEIL results****Figure 6: The clustering performance of the CEIL framework with GMM and CDCC over different iterations on various short text clustering tasks.**

that removing either the data filtering module or aggregation module will cause significant performance drop for GMM_CEIL. The reason is that in GMM_CEIL, the representation learning is only performed at the update of the classification task. Therefore, the quality of classification labels is critical for the clustering performance, which is greatly impacted by the data filtering and aggregation processes. In CDCC_CEIL, the performance decline of removing data filtering and aggregation modules is relatively small compared to GMM_CEIL, as the CDCC algorithm itself learns good text representations.

Prompt Learning. We adopt prompted input text segments for representation extraction, and prompt-tuning for the learning of the classifier. To demonstrate its effectiveness, we compare it with the normal fine-tuning approach, which uses the original text segments as the input for representation extraction and classification, and adopts an MLP network as the classification head for fine-tuning.

Results in Table 4 show that by using prompt learning, CDCC_p and CDCC_CEIL_p both achieve generally better clustering results compared with their fine-tuning counterparts. The success is contributed to the capability of prompt templates to better model the semantics when little supervision is provided.

Iterative Updating. In the CEIL framework, the clustering task and the classification task are supposed to promote each other in each iteration, and gain better clustering results and text representations iteratively. In order to verify this motivation, we conduct experiments on multiple short text clustering datasets and plot

the performance of GMM_CEIL and CDCC_CEIL over iterations, shown in Figure 6.

We can observe a rapid increase of clustering performance over iterations on GMM_CEIL shown in Figure 6a. This is because our CEIL framework can iteratively enhance the text representations based on the feedback of clustering results, leading to enormous improvements to traditional clustering algorithms like GMM which are highly dependent on the quality of data representations. A clear trend of performance increase over iterations can also be observed in Figure 6b, when CEIL is applied to powerful deep clustering algorithms like CDCC, which again illustrates the effectiveness and generality of our proposed CEIL framework.

5 CONCLUSION

In this paper, we propose a Classification-Enhanced Iterative Framework for short text clustering, where a classification objective is introduced in the clustering framework to iteratively enhance the text representations which further improves the clustering performance. We also propose a Category Disentangled Contrastive Clustering algorithm which outperforms other clustering algorithms by achieving better separation among clusters. Experimental results demonstrate that our CEIL generally improves the performance of various clustering algorithms and achieves the state-of-the-art performance on a wide range of short text clustering tasks.

REFERENCES

- [1] Sanjeev Arora, Yingyu Liang, and Tengyu Ma. 2017. A simple but tough-to-beat baseline for sentence embeddings. In *International conference on learning representations*.
- [2] Somnath Banerjee, Krishnan Ramanathan, and Ajay Gupta. 2007. Clustering short texts using wikipedia. In *Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval*. 787–788.
- [3] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*. PMLR, 1597–1607.
- [4] Cedric De Boom, Steven Van Canneyt, Thomas Demeester, and Bart Dhoedt. 2016. Representation learning for very short texts using weighted word embedding aggregation. *Pattern Recognition Letters* 80 (2016), 150–156.
- [5] 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* (2018).
- [6] Ning Ding, Shengding Hu, Weilin Zhao, Yulin Chen, Zhiyuan Liu, Hai-Tao Zheng, and Maosong Sun. 2021. Openprompt: An open-source framework for prompt-learning. *arXiv preprint arXiv:2111.01998* (2021).
- [7] Amir Hadifar, Lucas Sterckx, Thomas Demeester, and Chris Develder. 2019. A self-training approach for short text clustering. In *Proceedings of the 4th Workshop on Representation Learning for NLP (RepL4NLP-2019)*. 194–199.
- [8] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 9729–9738.
- [9] Xiaohua Hu, Xiaodan Zhang, Caimei Lu, Eun K Park, and Xiaohua Zhou. 2009. Exploiting wikipedia as external knowledge for document clustering. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*. 389–396.
- [10] Ryan Kiros, Yukun Zhu, Russ R Salakhutdinov, Richard Zemel, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. Skip-thought vectors. *Advances in neural information processing systems* 28 (2015).
- [11] Chen Li, Xiaoguang Yu, Shuangyong Song, Jia Wang, Bo Zou, and Xiaodong He. 2022. SimCTC: A Simple Contrast Learning Method of Text Clustering (Student Abstract). In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 36. 12997–12998.
- [12] Meilian Lu, Zhen Qin, Yiming Cao, Zhichao Liu, and Mengxing Wang. 2014. Scalable news recommendation using multi-dimensional similarity and Jaccard-Kmeans clustering. *Journal of Systems and Software* 95 (2014), 242–251.
- [13] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. *Advances in neural information processing systems* 26 (2013).
- [14] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. Automatic differentiation in pytorch. (2017).
- [15] Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*. 1532–1543.
- [16] Xuan-Hieu Phan, Le-Minh Nguyen, and Susumu Horiguchi. 2008. Learning to classify short and sparse text & web with hidden topics from large-scale data collections. In *Proceedings of the 17th international conference on World Wide Web*. 91–100.
- [17] Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. 2018. Improving language understanding by generative pre-training. (2018).
- [18] Md Rashadul Hasan Rakib, Norbert Zeh, Magdalena Jankowska, and Evangelos Milios. 2020. Enhancement of short text clustering by iterative classification. In *International Conference on Applications of Natural Language to Information Systems*. Springer, 105–117.
- [19] Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084* (2019).
- [20] Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of machine learning research* 9, 11 (2008).
- [21] Xiaojun Wan and Jianwu Yang. 2008. Multi-document summarization using cluster-based link analysis. In *Proceedings of the 31st annual international ACM SIGIR conference on Research and development in information retrieval*. 299–306.
- [22] Tingting Wei, Yonghe Lu, Huiyou Chang, Qiang Zhou, and Xianyu Bao. 2015. A semantic approach for text clustering using WordNet and lexical chains. *Expert Systems with Applications* 42, 4 (2015), 2264–2275.
- [23] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. 2018. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 3733–3742.
- [24] Junyuan Xie, Ross Girshick, and Ali Farhadi. 2016. Unsupervised deep embedding for clustering analysis. In *International conference on machine learning*. PMLR, 478–487.
- [25] Jiaming Xu, Bo Xu, Peng Wang, Suncong Zheng, Guanhua Tian, and Jun Zhao. 2017. Self-taught convolutional neural networks for short text clustering. *Neural Networks* 88 (2017), 22–31.
- [26] Bo Yang, Xiao Fu, Nicholas D Sidiropoulos, and Mingyi Hong. 2017. Towards k-means-friendly spaces: Simultaneous deep learning and clustering. In *international conference on machine learning*. PMLR, 3861–3870.
- [27] Jianwei Yang, Devi Parikh, and Dhruv Batra. 2016. Joint unsupervised learning of deep representations and image clusters. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 5147–5156.
- [28] Jianhua Yin and Jianyong Wang. 2016. A model-based approach for text clustering with outlier detection. In *2016 IEEE 32nd International Conference on Data Engineering (ICDE)*. IEEE, 625–636.
- [29] Dejiao Zhang, Feng Nan, Xiaokai Wei, Shangwen Li, Henghui Zhu, Kathleen McKeown, Ramesh Nallapati, Andrew Arnold, and Bing Xiang. 2021. Supporting clustering with contrastive learning. *arXiv preprint arXiv:2103.12953* (2021).
- [30] Xiang Zhang and Yann LeCun. 2015. Text understanding from scratch. *arXiv preprint arXiv:1502.01710* (2015).
- [31] Yi Zhang, Jie Lu, Feng Liu, Qian Liu, Alan Porter, Hongshu Chen, and Guangquan Zhang. 2018. Does deep learning help topic extraction? A kernel k-means clustering method with word embedding. *Journal of Informetrics* 12, 4 (2018), 1099–1117.
- [32] Li Zheng, Lei Li, Wenxing Hong, and Tao Li. 2013. PENETRATE: Personalized news recommendation using ensemble hierarchical clustering. *Expert Systems with Applications* 40, 6 (2013), 2127–2136.