


# Attentive Representation Learning with Adversarial Training for Short Text Clustering

Wei Zhang , Member, IEEE, Chao Dong, Jianhua Yin, and Jianyong Wang, Fellow, IEEE

**Abstract**—Short text clustering has far-reaching effects on semantic analysis, showing its importance for multiple applications such as corpus summarization and information retrieval. However, it inevitably encounters the severe sparsity of short text representations, making the previous clustering approaches still far from satisfactory. In this paper, we present a novel attentive representation learning model for short text clustering, wherein cluster-level attention is proposed to capture the correlations between text representations and cluster representations. Relying on this, the representation learning and clustering for short texts are seamlessly integrated into a unified model. To further ensure robust model training for short texts, we apply adversarial training to the unsupervised clustering setting, by injecting perturbations into the cluster representations. The model parameters and perturbations are optimized alternately through a minimax game. Extensive experiments on four real-world short text datasets demonstrate the superiority of the proposed model over several strong competitors, verifying that robust adversarial training yields substantial performance gains.

**Index Terms**—short text clustering, representation learning, attention mechanisms, robust adversarial training

## 1 INTRODUCTION

RECENT years have witnessed the fast-growing trade of short text data in various kinds of social media, for example, Twitter, Instagram, and Sina Weibo. As a consequence, short text clustering, the task of automatically grouping multiple unlabeled texts into a number of clusters, has become increasingly important. It can benefit multiple content-centric downstream applications, such as event exploration [12], trend detection [41], online user clustering [36], cluster-based retrieval [8], [37], to name a few. Compared with general text clustering, short text clustering is more challenging. This is because text representations in the original lexical space are usually sparse and this issue is further amplified for short texts [1]. Thus, the key to the success of short text clustering is to learn an effective short text representation scheme suitable for clustering.

On the basis of classical general clustering algorithms such as K-means [26], current developments in short text clustering mostly fall into two branches: Bayesian topic models (e.g., Latent Dirichlet Allocation (LDA) [6]) and deep learning approaches [33]. The former one realizes probabilistic text clustering by assuming that each document is associated with a distribution over topics, and each topic is a distribution over words. In this way, a topic is usually regarded as a cluster. To model short text, some elaborate topic models [62], [67] are presented by changing the text

generation process. However, one major limitation remains within most of these topic models: the input representation of short text is commonly based on the bag-of-words assumption and one-hot encoding, which might be sparse and lack the expressive ability.

Aiming to leverage the power of representation learning for short text clustering, the studies [60], [61] utilize word embeddings [42] and deep convolutional neural networks [28] to build a multi-stage framework. Better text representations are firstly learned to be fed into the conventional K-means algorithm for improving the clustering performance. However, the optimization process is partitioned into separate stages. Thus it is incapable of guiding text representation learning by the clustering objective. Some other deep representation learning approaches for general clustering problems [18], [27], [58], [69] have been trained in an end-to-end fashion, whereas they are not tailored for textual data or require an additional step to obtain short text representations. Specifically, hand-crafted text representations, such as Term Frequency-Inverse Document Frequency (TF-IDF), are commonly taken as model input. Hence they overlook the usage of word-level distributional representations, which are important to grasp the text semantic relatedness.

In this paper, we concentrate on bridging the gap between short text representation learning and short text clustering by fusing them in a unified model. Inspired by the roaring success of attention mechanisms [3] in Natural Language Processing (NLP), we devise the novel Attentive Representation Learning (ARL) model tailored for the short text clustering task. It first leverages low-dimensional word embeddings to build dense representations with the simple but effective mean-pooling technique for short texts. Cluster-level attention is then developed to capture the correlations between the short text representations and cluster representations, which shares a similar spirit with GSDMM [67] tailored for short text modeling that assumes all the words in a short text should belong to the same cluster (topic). The derived attention weights can be used as an evidence to determine cluster assignments. To enable the unsupervised learning of word and cluster representations, we propose to reconstruct short text representations through

- W. Zhang (corresponding author) and C. Dong are with the School of Computer Science and Technology, East China Normal University, Shanghai 200062, China (e-mail: zhangwei.thu2011@gmail.com; 51174500084@stu.ecnu.edu.cn).
- J. Yin is with Department of Computer Science and Technology, Shandong University, Qingdao 266237, Shandong, China (e-mail: jhyin@sdu.edu.cn).
- J. Wang is with the Department of Computer Science and Technology, Tsinghua University, Beijing 100086, China, and also with the Jiangsu Collaborative Innovation Center for Language Ability, Jiangsu Normal University, Xuzhou 221009, China (e-mail: jianyong@tsinghua.edu.cn).
- This paper was partially supported by the National Key Research and Development Program of China (No. 2020YFA0804503), the National Natural Science Foundation of China (No. 62072182, 61802231, 61532010, 61521002), and Beijing Academy of Artificial Intelligence (BAAI).

the weighted combination of cluster representations. An objective function combining the pairwise ranking loss and point-wise loss is employed to measure the reconstruction gap.

To improve the effectiveness of learning short text clustering, we further incorporate robust adversarial training [17], [43] into the final model, naming it as ARL-Adv. Robust adversarial training requires to feed both original real examples and intentional “adversarial examples” into the model during the training process. It has obtained impressive performance in the supervised and semi-supervised classification tasks. Concretely, ARL-Adv associates continuous cluster representations with adversarial perturbations to form an additional adversarial objective function. As a supplement to the original objective function, adversarial perturbations play a role of adaptive regularization, providing the optimization process with more robustness and being beneficial for short texts which sometimes contain noise. The model parameters and perturbations are optimized through a minimax game, where the model parameters are for minimizing the aforementioned reconstruction gap, and by contrast, the perturbations are for maximizing the reconstruction gap.

To sum up, our contributions are as follows:

- We present a novel attentive representation learning approach (ARL) to couple short text representation learning and clustering in a unified model. By proposing to reconstruct short-text representations through cluster-level attention, the model can be optimized in an unsupervised manner while cluster assignments are learned as well.
- To enhance the robustness of learning cluster label assignments for short texts, we introduce robust adversarial training by customizing adversarial examples (ARL-Adv). To our best knowledge, this is the first study to apply robust adversarial training into the field of the unsupervised clustering setting.
- We conduct extensive experiments on four real-world short text datasets. Three of them are publicly available, and the remaining one is constructed following a recent work. By comparing ARL-Adv with several strong baselines, we demonstrate its significant improvements and verify that the adversarial training further boosts performance. To ensure the reproducibility of this paper, we make the event-related dataset released for evaluation (see Section 5.4).

## 2 RELATED WORK

This section briefly reviews recent achievements on text clustering methods, especially for short texts, and generative adversarial networks for clustering. Besides, text representation schemes, attention mechanisms, and adversarial training are discussed.

### 2.1 Text Clustering Methods

While general text clustering has been extensively studied in the literature [1], [2], [23], [34], short text clustering [4], [48] receives less attention until recent years, due to the flourish of user-generated content in online social media. Currently, there are two main categories of approaches in the literature.

The first category is about topic modeling approaches, inheriting the ideas from LDA to uncover latent topic distributions by implicitly modeling word co-occurrence patterns. Unlike LDA that assumes to generate a topic for each word in a document, GSDMM [67] supposes all the words in a short text share the same topic, which is naturally regarded as the cluster of the text. Another

advanced model is BTM [62]. It explicitly models co-occurrence patterns through bigrams defined as pairs of unordered words in a given document. The cluster determination in BTM depends on the topic proportions of bigrams in a target short text. More recently, two topic models [35], [66] are developed for clustering streaming short texts incrementally. In this paper, we primarily concentrate on the general clustering setting and leave how to adapt our model to streaming short texts in the future.

Inspired by the success of deep learning [5], the latter category promotes its development in the clustering field. The most relevant study is STCC [60], [61] designed for short text clustering. However, it is not trained in an end-to-end fashion, leaving room for performance improvement. Gaussian-LDA [10] incorporates word embeddings into the generation process of topic models, but it is not very compatible with short texts, just as LDA. Recently, some deep learning approaches are proposed for general clustering tasks, but not crafted for texts. DEC [58] utilizes the idea of self-training to measure the gap between the distribution of soft cluster assignments and its corresponding auxiliary distribution through Kullback-Leibler (KL) divergence. STC [18] adopts self-training as well, meanwhile, Smoothed Inverse Frequency (SIF) is introduced to weight word embeddings in each document. Yet STC obtains text representations in a separate step and cannot optimize word embeddings along with learning text clustering. VaDE [27] extends variational auto-encoder [31] to integrate a Gaussian mixture model to generate latent vectors, instead of just using a Gaussian distribution. However, all the above models do not address the importance of representation learning for short texts, making their clustering performance not very satisfied in this situation.

### 2.2 Generative Adversarial Networks for Clustering

Benefiting from generative adversarial networks (GANs) [16] which aim to generate data that matches a target distribution, several approaches utilize GANs for clustering tasks. Specifically, Premachandran et al. [47] employed GANs to extract features, which are later fed into the K-means++ clustering algorithm. Therefore, it is not learned in an end-to-end fashion for clustering. DAC [19] leverages a discriminator to push forward the distribution of the data representations generated by a trainable encoder to a Gaussian mixture distribution. GANMM [69] revisits the expectation-maximization based clustering algorithm and replaces the expectation step and the maximization step with classification models and GANs. Mrabah et al. [45] extended DEC [58] by adding a regularization term through GANs to constrain the reconstructed data points to be similar to realistic data points. ClusterGAN [46] consists of a generator, a discriminator, and an encoder to learn mappings between discrete class labels and continuous image pixels. The innovation lies in the design of a discrete latent space  $Z$ , thereby providing sufficient signals for the generator to have samples belonging to a specific class and further benefiting clustering. The study [15] shares a similar spirit with ClusterGAN by devising a special latent space for ease of clustering. Moreover, self-paced learning is adopted to distinguish easy and hard examples for achieving a reliable training process. In summary, most of the GAN-based methods are welcomed for clustering images due to their ability to handle continuous values, whereas it is nontrivial for applying GAN-based methods to clustering the discrete textual data due to the difficulty of optimization [68]. A surrogate way is to represent text as fixed

continuous vectors (e.g., TF-IDF or pre-trained word embeddings) for GAN-based models. However, it might inevitably limit the power of learning text representations.

### 2.3 Text Representation Schemes

Traditionally, a single word is represented by a one-hot vector with a fixed dimensional space. The space has the same size as that of a specified vocabulary. Accompanied by this, vector space models with TF-IDF [40] are commonly utilized to denote short text representations. To alleviate the issue of sparse exact word overlap in similarity computation, low-dimensional techniques, such as Principal Component Analysis (PCA) and probabilistic topic models (e.g., LDA), are applied to the text domain. In particular, word embedding techniques [42] have gained lots of attention in recent years, due to their improved ability to grasp semantic relatedness. Consequently, much progress has been observed in this respect, including contextualized word embeddings like BERT [11] and GPT-3 [7]. Benefiting from the advantage of word embeddings, the pre-training and fine-tuning paradigm [59] becomes the state-of-the-art for text classification and some other NLP tasks. Nevertheless, the embedding based text clustering approaches [18], [58], [60] could not fine-tune the fine-grained word-level representations by the guidance of clustering objectives, thus leaving room for improving text representations. This motivates the proposed models ARL and ARL-Adv to mitigate this issue by jointly learning word embeddings and text clustering (see the degraded performance of not fine-tuning word embeddings in Table 4).

### 2.4 Attention Mechanisms

Attention mechanisms are growing in popularity for learning representations for image, text, and other different modalities [3], [14], [54], [57], [65]. For a given context (query), it typically calculates the prominence of each representation (key) in sub-layers to form an integrated representation (value). In the literature of text modeling, attention mechanisms allow for dependency modeling regardless of the distance between words, compared with recurrent neural networks [22]. The recent study [54] develops a multi-header self-attention mechanism for representing input and output sequences. While the approach is entirely composed of attention blocks, it outperforms several recurrence and convolution based models in machine translation tasks. Based on this self-attention mechanism, large-scale pre-trained language models (e.g., BERT) have been proposed and benefited many NLP tasks. For the specific text classification task, Ma et al. [38] introduced a mutual attention mechanism to assign importance weights to both semantic related long-distance words and local semantic features. In another aspect, the works [9], [55] consider incorporating user factors into the attention computation process for words and sentences. The work [20] also adopts word-level attention to extract aspect-related words and de-emphasize irrelevant words. Different from these studies, we perform cluster-level attention to associate the representation learning with the automatic determination of cluster assignments in a unified model to achieve the state-of-the-art short text clustering performance.

### 2.5 Robust Adversarial Training

Since the pioneering work [17] proposes to investigate the benefit of supervised learning on intentional adversarial examples for image classification, robust adversarial training has made continuous

progress in different fields [21], [43], [44], [63]. For example, Miyato et al. [43] applied virtual adversarial training [44], a variant of adversarial training for semi-supervised learning, to the text classification task. He et al. [21] learned a pairwise ranking based matrix factorization model for achieving better recommendation performance in an adversarial manner. In a nutshell, the peculiarities of robust adversarial training are to proactively inject small intentional perturbations to the original data and train models to behave well on them, thereby enhancing the robustness of model training. To the best of our knowledge, this is the first study to leverage robust adversarial training in the literature of unsupervised clustering.

It is worth noting that robust adversarial training has a remarkable difference with GANs due to the following reasons. First of all, unlike GANs aiming at generating new data to match a specified target distribution, robust adversarial training targets at ensuring a robust model training process. Secondly, there is no need for robust adversarial training to have a tailored discriminator like GANs since the target is not to distinguish real or synthetic data. Because of this, some studies [13], [70] that couple GANs with attention mechanisms do not share the same technological thought, let alone they are not proposed for clustering tasks. Besides, although Wei et al. [56] adopted adversarial attacks for clustering, their approach is largely different from the proposed ARL-Adv since: (1) It relies on a few labeled data in its clustering procedure. (2) The grid based clustering and adversarial attack are separated in different stages. (3) The adversarial part does not involve model parameter learning.

## 3 THE COMPUTATIONAL MODEL

**Model Overview:** Figure 1 depicts the overall architecture of the proposed model ARL-Adv with given short text examples. From a whole perspective, the model contains five computational modules: (1) short text representation, (2) cluster-level attention, (3) short text reconstruction, (4) objective function, and (5) robust adversarial training. Taking the given short text (i.e., text  $i$ ) as an illustrative example, it is first fed into the short text representation module (see Section 3.1), where the words (e.g., tax) in the text are associated with trainable word embeddings to constitute its text representation  $d_i$ . Afterwards, cluster-level attention (see Section 3.2) is conducted to compute the probability (relevance) weights between the text representation and different cluster representations (e.g.,  $c_1$ ) that are trainable as well. The weights are consequently adopted to combine cluster representations for reconstructing short text representation  $\tilde{d}_i$  (see Section 3.3). To build the pairwise ranking based clustering objective function (see Section 3.4), short text  $j$  is sampled from the given corpus as a pseudo negative text. Eventually, the adversarial training module (see Section 4) injects perturbations (e.g.,  $\Delta_1^C$ ) into the cluster representations to constitute another short text reconstruction  $\tilde{d}_i^\Delta$  for ranking, thereby providing the ability of robust training. After ARL-Adv is well trained, the cluster label of short text  $i$  is determined by sorting the corresponding probability weights and returning the cluster index (e.g.,  $k$ ) with the maximal weight (e.g.,  $p_{i,k}$ ).

### 3.1 Short Text Representation

For the short text clustering problem, we have a corpus  $\mathcal{D}$  containing multiple short texts, i.e.,  $\{D_1, D_2, \dots, D_{|\mathcal{D}|}\}$ , where

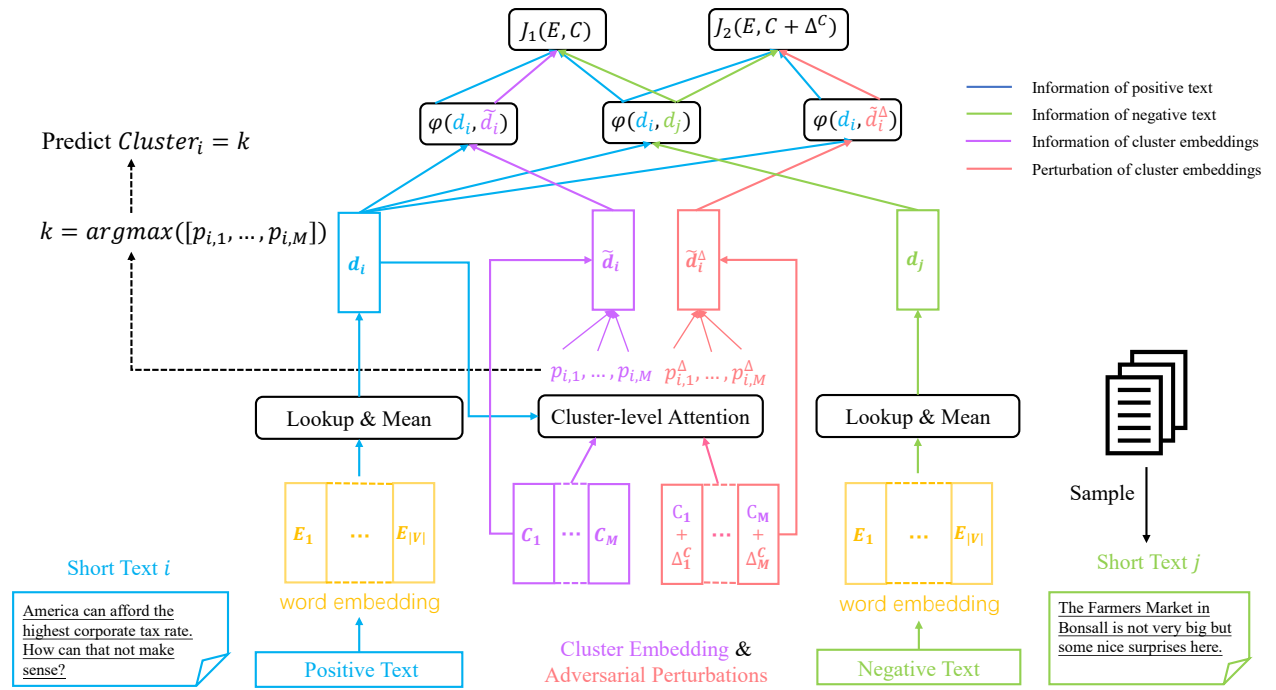


Fig. 1: Architecture of the proposed ARL-Adv. Four types of marks with different colors are used to label different information flow. The dotted lines in the middle left of the figure denote how the cluster label of short text  $i$  is derived after the model is well trained.  $J_1$  and  $J_2$  correspond to the normal objective function (Equation (8)) and the adversarial objective function (Equation (9)), respectively.

$|\mathcal{D}|$  is the size of  $\mathcal{D}$ . Taking the  $i$ -th short text  $D_i \in \mathcal{D}$  for illustration, it is associated with a matrix based representation  $\mathbf{D}_i = \{\mathbf{x}_t^i\}_{t=1}^{l_i}$ , where  $l_i$  is the length of the text and  $\mathbf{x}_t^i$  is the column vector of one-hot encoding at position  $t$ , with the size equal to the number of words in a pre-specified vocabulary  $V$ . It is worth noting that here  $\mathbf{D}_i$  is the original text representation, while some other strategies such as TF-IDF can be regarded as the hand-crafted feature engineering for texts.

To convert each one-hot sparse representation  $\mathbf{x}_t^i$  ( $t \in \{1, \dots, l_i\}$ ) to its low-dimensional dense embedding, we first adopt a look-up encoding procedure:

$$\mathbf{w}_t^i = \mathbf{E}\mathbf{x}_t^i, \quad (1)$$

where  $\mathbf{E} \in \mathbb{R}^{K \times |V|}$  is a trainable word embedding matrix and  $K$  is the embedding size. Given these dense word embeddings, a simple mean-pooling technique is empirically adopted to construct an informative representation of the short text:

$$\mathbf{d}_i = \frac{1}{l_i} \sum_{t=1}^{l_i} \mathbf{w}_t^i. \quad (2)$$

The motivation behind the above operation is that short texts commonly have insufficient contextual information and might sometimes contain noisy word information. As such, complex contextual text representation methods, including pre-trained contextual word representations (e.g., BERT), attention based document-level representations [64], recurrent sequential representations [53], and convolution-based representations [29], do not exhibit significant improvement and are sometime even worse than mean-pooling in our local tests. We have also attempted VLAWE [25] for text representations, whereas it is not very suitable because the word embeddings are hard to be fine-tuned

when learning short text clustering due to its separate procedure of word embedding clustering. Consequently, we deem that mean-pooling is suitable for the simultaneous learning of short text representations and clustering.

### 3.2 Cluster-level Attention

To represent the presumed  $M$  clusters for the specific corpus  $\mathcal{D}$ , a trainable cluster representation matrix  $\mathbf{C} = [\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_M] \in \mathbb{R}^{K \times M}$  is established. As  $\mathbf{d}_i$  has contained useful global semantic information for the  $i$ -th short text, leveraging it to select a suitable cluster can be feasible, which is the primary motivation of cluster-level attention. In general, an attention function takes a query and a set of key-value pairs as input. Each attention score characterizes the compatibility of its corresponding key with the query [54], and is later used for the weighted combination of all values. We let both key and value correspond to the representations of clusters, and regard short text representations as queries. It is intuitive to perform selection based on a probability distribution  $p(z_m | \mathbf{d}_i)$  over each cluster  $z_m$  where  $m \in \{1, 2, \dots, M\}$ . Formally, the distribution is defined as follows:

$$z_m \sim p(z_m | \mathbf{d}_i) = \frac{\exp(\mathbf{c}_m^\top \mathbf{d}_i)}{\sum_{m'=1}^M \exp(\mathbf{c}_{m'}^\top \mathbf{d}_i)}. \quad (3)$$

The above computational manner is the realization of cluster-level attention, which captures the text-cluster interaction. The idea behind it also conforms to the successful experience of probabilistic topic models for short text modeling. Specifically, in LDA-like models, the topic of each word in a long text is separately determined from their sampled topics. By contrast, GSDMM assumes that all the words in the same short text belong to the same topic. It exhibits improved performance in short



text modeling. ARL-Adv is consistent with GSDMM, which has only one latent topic  $z_m$  for a whole short text, as shown in Equation (3).

We denote  $p(z_m|\mathbf{d}_i)$  as  $p_{i,z_m}$  for short. By referring to the probability values  $p_{i,z_m}$  for different  $m$ , it is sufficient to identify the most relevant cluster given  $\mathbf{d}_i$ . Besides, each probability value could be employed to rescale the representations of the corresponding clusters for  $\mathbf{d}_i$ , yielding  $\hat{\mathbf{C}}_i = [\hat{\mathbf{c}}_{i,1}, \hat{\mathbf{c}}_{i,2}, \dots, \hat{\mathbf{c}}_{i,M}]$ , where  $\hat{\mathbf{c}}_{i,m} = p_{i,z_m} \mathbf{c}_m$ .

### 3.3 Short Text Reconstruction

Since short text clustering is indeed an unsupervised learning problem, we leverage the cluster representations to reconstruct the short text representation, which is later utilized to guide the training of the whole model.

By assuming the reconstructed representation of the  $i$ -th text as  $\tilde{\mathbf{d}}_i$ , we calculate it based on a linear combination of the target-dependent cluster matrix  $\hat{\mathbf{C}}_i$ :

$$\tilde{\mathbf{d}}_i = \sum_{m=1}^M \hat{\mathbf{c}}_{i,m}. \quad (4)$$

Benefiting from the different contributions of clusters, we can rebuild an expressive representation of the text. For a specific text, the derivation of cluster-level attention in Equation (3) encourages those clusters that are more similar to the text representation to contribute more in the linear combination in Equation (4). When the derived reconstruction is driven to be closer to the text representation, cluster-level attention may gradually favor one single cluster as the training goes on. The visualization in Figure 2 can partly explain the consequence. In the experiments, we find the average probabilities of the inferred maximum cluster are over 0.99. These results mean each text is highly concentrated to a single cluster, rather than an arbitrary value.

### 3.4 Objective Function

Given the obtained short text representation  $\mathbf{d}_i$  and its reconstructed representation  $\tilde{\mathbf{d}}_i$  from cluster representations, we define their relevance score to be the cosine similarity between the two representations:

$$\varphi(\mathbf{d}_i, \tilde{\mathbf{d}}_i) = \frac{\mathbf{d}_i^\top \tilde{\mathbf{d}}_i}{\|\mathbf{d}_i\| \|\tilde{\mathbf{d}}_i\|}. \quad (5)$$

We deem that when the relevance score gets larger, the result of the reconstruction is better. To achieve this, we propose a hybrid objective function, consisting of two parts: a pairwise ranking part and a relevance maximization part.

As for the first part, our model seeks to minimize a margin based pairwise ranking loss [51]. Specifically, for the  $i$ -th short text, we sample several other texts from the corpus and regard them as a pseudo negative text set  $\mathcal{N}_i$ . And for text  $D_j \in \mathcal{N}_i$ , we obtain the text similarity of the two short texts, i.e.,  $\varphi(\mathbf{d}_i, \mathbf{d}_j)$ , based on the cosine similarity of their text representations calculated by Equation (5), where embedding  $\mathbf{d}_j$  is obtained with the same procedure as shown in Equation (2). Afterwards, we formulate the pairwise ranking loss for the  $i$ -th short text as follows:

$$\mathcal{L}_1(i; \mathbf{E}, \mathbf{C}) = \frac{1}{|\mathcal{N}_i|} \sum_{j \in \mathcal{N}_i} \max(0, \gamma - \varphi(\mathbf{d}_i, \tilde{\mathbf{d}}_i) + \varphi(\mathbf{d}_i, \mathbf{d}_j)), \quad (6)$$

where  $|\mathcal{N}_i|$  denotes the number of pseudo negative text (empirically setting to 16), and  $\gamma$  is the margin between positive and negative ones (empirically setting to 1). We can see minimizing the above loss would make the reconstructed  $\tilde{\mathbf{d}}_i$  to have a larger relevance score with the original representation  $\mathbf{d}_i$  than with negative samples at the given margin.

As a supplement, we adopt the relevance maximization of  $\varphi(\mathbf{d}_i, \tilde{\mathbf{d}}_i)$ . To keep consistent with the optimization direction of  $\mathcal{L}_1$ , we define the following loss function:

$$\mathcal{L}_2(i; \mathbf{E}, \mathbf{C}) = -\varphi(\mathbf{d}_i, \tilde{\mathbf{d}}_i), \quad (7)$$

which can be regarded as a pointwise loss function, aiming to predict  $\mathbf{d}_i$  as similar as possible given the combination of all cluster representations. By fusing  $\mathcal{L}_1(i)$  and  $\mathcal{L}_2(i)$ , it ensures that  $\varphi(\mathbf{d}_i, \tilde{\mathbf{d}}_i)$  not only has larger relevance than those of pseudo negative text pairs, but also a large value itself. Taking all the short texts in corpus  $\mathcal{D}$  into consideration, we define the overall objective loss function as follows:

$$J_1(\mathbf{E}, \mathbf{C}) = \sum_{i=1}^{|\mathcal{D}|} (\mathcal{L}_1(i; \mathbf{E}, \mathbf{C}) + \mathcal{L}_2(i; \mathbf{E}, \mathbf{C})). \quad (8)$$

From a whole perspective, ARL-Adv contains word embedding matrix  $\mathbf{E}$  and cluster representation matrix  $\mathbf{C}$  as its parameters. Each short text is directly encoded into a low-dimensional space through its word embeddings so that no extraordinary parameters are required.

## 4 ADVERSARIAL TRAINING FOR CLUSTERING

We propose to apply robust adversarial training to facilitate the representation learning process of ARL-Adv. In essence, robust adversarial training provides a new objective function based on adversarial perturbations to complement the original optimization procedure. This objective requires the model to perform well on both original samples and adversarial samples, thus benefiting the model in its training robustness, especially for short texts.

Since typical adversarial perturbations take continuous values and are added to real-valued vectors like images, they are not directly applied to discrete tokens like words. Following the study [44], we add continuous adversarial perturbations to the cluster representations. It is worth noting that we do not add adversarial perturbations to the word embedding matrix. The reasons lie in the two aspects, i.e., efficiency and effectiveness. First, the number of words extends far beyond that of clusters. Thus adding perturbations to words instead of clusters entails many more parameters (perturbations) to be optimized, which might reduce the model training efficiency. Second, we empirically validate that applying perturbations to words is less effective than adding perturbations to cluster representations, the results of which are shown in Table 4. Formally, we define the adversarial cluster-level perturbations in ARL-Adv as  $\Delta^C \in \mathbb{R}^{K \times M}$ . Consequently, we propose a new objective function formulated as follows:

$$J_2(\mathbf{E}, \mathbf{C} + \Delta^C) = \sum_{i=1}^{|\mathcal{D}|} (\mathcal{L}_1(i; \mathbf{E}, \mathbf{C} + \Delta^C) + \mathcal{L}_2(i; \mathbf{E}, \mathbf{C} + \Delta^C)), \quad (9)$$

where the new cluster representations  $\mathbf{C} + \Delta^C$ , accompanied by the corresponding probability distribution  $p_{i,z_m}^\Delta$  where  $m \in \{1, 2, \dots, M\}$ , are leveraged to constitute another reconstructed

representation  $\tilde{\mathbf{d}}_i^\Delta$  for short text  $\mathcal{D}_i$ . The optimization of the above objective function can be viewed from two aspects. For learning  $\Delta^C$ , the optimal condition is as follows:

$$\Delta^{C*} = \arg \max_{\Delta^C} J_2(\mathbf{E}, \mathbf{C} + \Delta^C). \quad (10)$$

On the contrary, the optimization target of  $\mathbf{E}$  and  $\mathbf{C}$  is to minimize  $J_2$ . Typically, a norm-based constraint should be adopted to restrict the scale of  $\Delta^C$ .

By integrating the two objective functions, i.e.,  $J_1$  and  $J_2$ , the final optimization target is given as:

$$J(\mathbf{E}, \mathbf{C}, \Delta^C) = J_1(\mathbf{E}, \mathbf{C}) + \alpha J_2(\mathbf{E}, \mathbf{C} + \Delta^C), \quad (11)$$

$$\mathbf{E}^*, \mathbf{C}^*, \Delta^{C*} = \arg \min_{\mathbf{E}, \mathbf{C}} \max_{\Delta^C} J(\mathbf{E}, \mathbf{C}, \Delta^C), \quad (12)$$

where  $\alpha$  controls the relative strength of  $J_2$ . The optimization of Equation (11) involves playing a minimax game. At each step, the worst-case perturbations to cluster representations are first identified by increasing the value of  $J_2$  as much as possible. Afterwards,  $\mathbf{E}$  and  $\mathbf{C}$  are optimized to be robust to such intentional perturbations in  $J_2$ , while retaining satisfactory performance in  $J_1$ . In such a way, for text representation  $\mathbf{d}_i$ , its correlations with the normal reconstruction  $\tilde{\mathbf{d}}_i$  and the adversarial reconstruction  $\tilde{\mathbf{d}}_i^\Delta$  are simultaneously considered. This ensures the robustness of model training, which conforms to the short-length and noisy situation of short texts.

Algorithm 1 summarizes the whole optimization process. As usual, stochastic gradient descent algorithms such as Adam [30] are leveraged to learn the word and clustering embeddings. ARL-Adv calculates the gradients of  $\mathbf{E}$  and  $\mathbf{C}$  over Equation (11) through back propagation, and updates the parameters accordingly. The learning of adversarial perturbations  $\Delta^C$  follows the approximated linearizing methodology [17], leading to fast solutions with analytic forms. With the adding of L2 norm to each column of  $\Delta^C$ , i.e.,  $\|\Delta_m^C\| \leq \epsilon$  ( $m \in \{1, 2, \dots, M\}$ ), we can easily derive the following updating rule for the perturbation of each cluster representation:

$$\Delta_m^C = \epsilon \frac{\mathbf{g}_m^C}{\|\mathbf{g}_m^C\|}, \quad \text{where } \mathbf{g}_m^C = \alpha \frac{\partial J_2}{\partial \mathbf{C}_m}. \quad (13)$$

After finishing the training of ALR-Adv, the cluster assignments are determined based on the probability weights over the learned cluster embedding matrix  $\mathbf{C}$ , while the adversarial perturbations  $\Delta^C$  are ignored.

## 5 EXPERIMENTS

In this section, we conduct experiments to answer the following pivotal questions:

- Q1 Does the proposed ARL-Adv outperform standard and advanced text clustering models on short text datasets?
- Q2 Are the main components of the proposed model, including robust adversarial training, beneficial for the performance of short text clustering?
- Q3 How good are the learned short text embeddings and cluster representations, compared to other baselines?

### 5.1 Datasets

We utilize four real-world short text datasets. A brief introduction of each dataset and some common text preprocessing procedures are provided in the following.

---

#### Algorithm 1: Robust adversarial training for ARL-Adv.

---

**Input** : Short text corpus  $\mathcal{D}$ , cluster number  $M$ , adversarial strength  $\alpha$ , adversarial norm  $\epsilon$ , and some other hyper-parameters such as embedding size  $K$ .

**Output**: Word embeddings  $\mathbf{E}$ , cluster representations  $\mathbf{C}$ , cluster-level attention weights  $p(z_m|\mathbf{d}_i)$  ( $\forall m \in \{1, 2, \dots, M\}$  and  $i \in \{1, 2, \dots, |\mathcal{D}|\}$ ).

```

1 begin
2   Initialize word embeddings  $\mathbf{E}$  and cluster
   representations  $\mathbf{C}$ ;
3   for  $iter = 1$  to  $max\_iter$  do
4     Sample a mini-batch  $\mathcal{D}_{batch} \in \mathcal{D}$ ;
5     for  $i \in \mathcal{D}_{batch}$  do
6       Constructing text pairs  $(i, j)$ , where
          $j \in \mathcal{D}_{batch}$ ;
7       Learning adversarial perturbations  $\Delta^C$ 
         through Equation (13);
8       Updating word embeddings  $\mathbf{E}$  and cluster
         representations  $\mathbf{C}$  by gradient descent of
         Equation (12);

```

---

- **TREC**<sup>1</sup>: The dataset is from the Text REtrieval Conference on 2011-2015 tweet tracks. They are organized by their corresponding queries and evaluated into several relevance levels. We retain tweets labeled relevant or highly-relevant to their queries to ensure the quality of labels.
- **GoogleNews**<sup>1</sup>: This dataset is composed of groups of news titles and snippets clawed from Google News<sup>2</sup>. Manual observation has confirmed its favorable grouping quality.
- **Event**: Following the study [49], we extract event-related tweets from an off-the-shelf tweet dataset crawled in 2016<sup>3</sup>. Prior knowledge about the events, including the time window, relevant entities, and keywords, is fetched from Wikipedia.
- **StackOverflow**<sup>4</sup>: It is created based on the questions posted in Stack Overflow. We require each of the selected questions to be associated with only one tag. And the tags are regarded as the ground-truth cluster labels for the questions. The created dataset is substantially larger than the above three datasets, hoping to provide a more convincing empirical study. For simplicity, we sometimes use the name SO to denote this dataset.

For preprocessing, we utilize typical procedures in text processing, which consists of converting words into lowercase, removing stop words and irregular words, and applying Porter stemming. In addition, words with frequencies below 5 are discarded. The detailed statistics of all datasets can be found in Table 1. *#Cluster* and *#Text* refer to the actual number of clusters and texts for each dataset, respectively. *Vocabulary* is the number of remaining word tokens after preprocessing. *Avg. Len.* stands for the average length of texts counted in words.

1. <https://github.com/jackyin12/MStream>
2. <http://news.google.com>
3. <https://archive.org/>
4. <https://archive.org/download/stackexchange>

TABLE 1: Detailed statistics of each dataset.

Dataset	#Cluster	#Text	Vocabulary	Avg. Len.
TREC	128	11142	3503	7.87
GoogleNews	152	11109	6555	6.23
Event	69	26619	3314	8.78
StackOverflow	111	137367	10832	5.42

TABLE 2: Default settings for ARL-Adv.

$K$	$ \mathcal{D}_{batch} $	$\alpha$	$\epsilon$
300	64	1.0	50.0

## 5.2 Methods for Comparisons

### 5.2.1 Baselines

The selected representative baselines are categorized into conventional text clustering methods (i), deep learning based general clustering models (ii), and tailored methods for short text clustering (iii). The details of the baselines are given below for clarity.

- **K-means**: It is a classical and simple clustering approach, relying on hand-crafted features for text clustering. Two feature types are adopted in our experiments, i.e., TF-IDF and low-dimensional representations obtained by PCA. We denote K-means running on these two different features as K-means(TF-IDF) and K-means(PCA), respectively. For K-means(PCA), we have tried the number of primary components in  $\{100, 200, 300\}$ . It turns out 200 is the more suitable choice, so we report results based on this setting. The Euclidean distance is utilized as the internal similarity metric. For each run, 10 random centroids are initialized.
- **HieClu**: Similar as K-means, Hierarchical Clustering (HieClu) is another simple baseline for clustering. The results of taking textual features from PCA are reported due to its better performance than TF-IDF.
- **LDA** [6]: LDA is a classical and standard generative statistical model which learns a topic (cluster) distribution for each document. We assign a text with the topic that has the largest probability value inferred by LDA. Following [50], we set  $\alpha = K/50$  and  $\beta = 0.1$ , where  $K$  is the number of topics. Other parameter settings are tried as well, but we observe no significant improvement in performance.
- **BTM** [62]<sup>5</sup>: BTM regards bi-grams as the representations of short texts and generates them conditioned on different topics. Similar to LDA, we tune  $\alpha = K/50$  and set  $\beta = 0.05$  based on performance. The number of training iterations is set to 500 for BTM to reach convergence.
- **GSDMM** [67]<sup>6</sup>: GSDMM is tailored for short text clustering and assumes that each text is generated by one topic, which is fundamentally different from LDA that generates one word from one topic. We set the parameters  $\alpha = K/50$  and  $\beta = 0.05$ , and train the model for 300 iterations.
- **DEC** [58]<sup>7</sup>: It is a deep embedded clustering model that leverages autoencoder, with TF-IDF features as input to map documents into low-dimensional embeddings. Then the mapping function and cluster representations are refined based on the idea of self-training. Similar to VaDE, DEC has specified transforming and parameter settings for text data. We perform pre-training for 500 iterations, and training for 10000 iterations to ensure its convergence. SGD is the optimizer during both pre-training and training phases, with momentum set to 0.9. The batch size is set to 128. Empirical results show that setting

the dimension of the hidden layer in DEC to 80 can help it get consistent and robust performance.

- **STCC** [60], [61]<sup>8</sup>: This model mainly consists of three separate steps. It first trains a convolutional neural network with the help of autoencoders. Afterwards, the well-trained model is employed to get text embeddings, which are fed into K-means for final clustering. We set the kernel width to 3 for convolutional layers, with  $k = 5$  in its  $k$ -max pooling. The number of feature maps on the first and second convolutional layer is 32 and 16, respectively. Adam is used for batch based optimization, with the learning rate of  $5e-3$  and batch size of 100.
- **VaDE** [27]<sup>9</sup>: VaDE extends canonical variational auto-encoding approaches to support clustering tasks by utilizing the idea of Gaussian mixture model. We partially follow the parameter settings for text data in the implementation of VaDE and keep the procedure of TF-IDF based feature transformation. Setting pre-training to 500 iterations and training to 10000 iterations results in the convergence of VaDE on an empirical basis. Both phases use Adam as the optimizer, with the learning rate of  $1e-3$ . The default configuration of Adam is adopted, i.e.,  $\beta_1 = 0.9$  and  $\beta_2 = 0.999$ . We adjust the size of the hidden representations to 40, and batch size to 128 for more general and stable outcomes.
- **ClusterGAN** [46]<sup>10</sup>: ClusterGAN is recently proposed to leverage GANs to encode continuous data to discrete cluster labels. Considering its requirement of continuous data representations, we have tried TF-IDF and word2vec (average word embedding) to denote each short text, and reported the better performance.
- **STC** [18]<sup>11</sup>: STC adopts self-training inspired by DEC and uses Smoothed Inverse Frequency (SIF) to compute a weighted average of pre-trained word embeddings in a stage independent of optimizing its clustering model. The word embeddings are initialized by word2vec (the same as ARL-Adv), and their size is tuned to be 32. The dimension of its autoencoder is set to 80. Other parameters are kept the same as their original settings.

All the above models are tuned for the experimental datasets. To ensure statistical significance, all results in the experiments are averaged over ten runs. We also tried GaussianLDA<sup>12</sup> and GANMM<sup>13</sup> introduced in Section 2.1. However, the results show that GaussianLDA could not obtain competitive clustering performance than other baselines. And training GANMM incurs a heavy computational burden, especially when the cluster number is not small. Therefore we omit their details in this paper.

### 5.2.2 Variants of the proposed model

To verify the effectiveness of the main components in the proposed model, we consider some of its variants, which are introduced below.

8. <https://github.com/jacoxu/STC2>
9. <https://github.com/slim1017/VaDE>
10. <https://github.com/sudiptodip15/ClusterGAN>
11. [https://github.com/hadifar/stc\\_clustering](https://github.com/hadifar/stc_clustering)
12. [https://github.com/rajarshd/Gaussian\\_LDA](https://github.com/rajarshd/Gaussian_LDA)
13. <https://github.com/eyounx/GANMM>

5. <https://github.com/xiaohuiyan/BTM>  
 6. <https://github.com/jackyin12/GSDMM>  
 7. <https://github.com/piiswong/dec>

TABLE 3: Performance comparison over all methods.

Type	Dataset Metrics	TREC			GoogleNews			Event			StackOverflow		
		NMI	ARI	ACC	NMI	ARI	ACC	NMI	ARI	ACC	NMI	ARI	ACC
i	LDA	0.7514	0.5897	0.6120	0.7197	0.5074	0.5878	0.6348	0.4557	0.4984	0.2121	0.1044	0.2049
	K-means(TF-IDF)	0.8376	0.3682	0.6463	0.7941	0.2316	0.5897	0.6790	0.3290	0.4473	0.3110	0.0682	0.2459
	K-means(PCA)	0.8289	0.3460	0.6267	0.7724	0.2011	0.5547	0.6680	0.3150	0.4332	0.2885	0.0848	0.2319
	HieClu	0.8567	0.4975	0.7118	0.7816	0.2459	0.5897	0.5901	0.2346	0.4013	0.2756	0.0336	0.2229
ii	ClusterGAN	0.6445	0.3602	0.4953	0.7206	0.4058	0.5396	0.4061	0.1211	0.2247	0.2217	0.0844	0.2038
	DEC	0.8664	0.5839	0.6993	0.8505	0.5088	0.6705	0.7423	0.4264	0.5320	0.3799	0.1516	0.3186
	VaDE	0.8712	0.6632	0.7170	0.8389	0.5318	0.6287	0.6620	0.2927	0.4078	0.3214	0.1277	0.2670
iii	STCC	0.8172	0.5758	0.6455	0.7947	0.5262	0.6122	0.6567	0.3797	0.3959	0.3608	0.1433	0.2522
	BTM	0.8759	0.6920	0.7182	0.8656	0.6310	0.7023	0.6781	0.3834	0.4704	0.3508	0.1537	0.2565
	GSDMM	0.8746	0.7453	0.7512	0.8700	0.6782	0.7278	0.7572	0.4937	0.5039	0.3569	0.1418	0.2598
	STC	0.8906	0.7372	0.7589	0.8667	0.6471	0.7169	0.6431	0.3170	0.3982	0.3903	0.1598	0.2750
	<b>ARL</b>	0.9261	0.8428	0.8347	0.8995	0.7897	0.8057	0.8243	0.6078	0.6156	0.4796	0.2731	0.3781
	<b>ARL-Adv</b>	<b>0.9305</b>	<b>0.8486</b>	<b>0.8402</b>	<b>0.9054</b>	<b>0.8214</b>	<b>0.8303</b>	<b>0.8450</b>	<b>0.6728</b>	<b>0.6607</b>	<b>0.4871</b>	<b>0.3006</b>	<b>0.3943</b>

- **ARL-Adv**: This is the full version of our proposed model, which adopts the optimization target as Equation (11) and enables adversarial training.
- **ARL**: Only  $J_1$  is employed to train the model. By comparing ARL with the full model ARL-Adv, we can verify the effectiveness of robust adversarial training for the unsupervised clustering task.
- **ARL-Adv(no train w)**: This method optimizes parameters other than the word embedding matrix. Comparison against it can show the necessity of learning word-level semantic representations.
- **ARL-Adv(no train c)**: The cluster representation matrix is not learned in this model, hoping to partially verify the benefit of combining representation learning and clustering.
- **ARL-Adv w/o  $\mathcal{L}_1$** : The pairwise ranking loss  $\mathcal{L}_1$  is removed from ARL-Adv. That is, the short texts do not interact with their negative samples.
- **ARL-Adv w/o  $\mathcal{L}_2$** : It means ARL-Adv does not consider the supplementary pointwise loss  $\mathcal{L}_2$ . Both ARL-Adv w/o  $\mathcal{L}_1$  and ARL-Adv w/o  $\mathcal{L}_2$  still perform adversarial training.
- **ARL-Random**: This variant takes random noise as perturbations, with the same scale level as intentional adversarial perturbations.
- **ARL-Adv(word)**: It shares a very similar spirit with ARL-Adv, except that adversarial perturbations are added to word embeddings instead of cluster representations.

For ARL-Adv and its variants, word embeddings are pre-trained on each dataset with word2vec<sup>14</sup> for 200 iterations. Moreover, we initialize centroids of the clusters by performing K-means on short text embeddings, the same as the study [58]. Without loss of generality, the default hyper-parameters are set to the ones shown in Table 2.  $K$  is the dimension of embeddings, and  $|\mathcal{D}_{batch}|$  is the size of a batch. Unless otherwise stated, the results will be reported under this setting.

### 5.3 Evaluation Metrics

We adopt three commonly used metrics for text clustering performance evaluation, i.e., normalized mutual information (NMI) [52], adjusted rand index (ARI) [24], and clustering accuracy (ACC) [58]. Suppose  $\nu_i$  denotes the number of short texts in the  $i$ -th true topic, and  $\tilde{\nu}_j$  represents the number of short texts in the  $j$ -th inferred cluster. We use  $\bar{\nu}_{ij}$  to denote the number of short texts simultaneously appearing in two clusters. In addition,

$\mathcal{M}$  denotes the set of all possible one-to-one mappings between the generated clusters and real topics. For an efficient search of the best mapping, Hungarian algorithm [32] can be adopted. Formally, NMI, ARI, and ACC can be formulated as follows:

$$NMI = \frac{\sum_{i,j} \bar{\nu}_{ij} \log \frac{|\mathcal{D}| \cdot \bar{\nu}_{ij}}{\nu_i \tilde{\nu}_j}}{\sqrt{(\sum_i \nu_i \log \frac{\nu_i}{|\mathcal{D}|}) (\sum_j \tilde{\nu}_j \log \frac{\tilde{\nu}_j}{|\mathcal{D}|)}}, \quad (14)$$

$$ARI = \frac{\sum_{i,j} \binom{\bar{\nu}_{ij}}{2} - \left[ \sum_i \binom{\nu_i}{2} \sum_j \binom{\tilde{\nu}_j}{2} \right] / \binom{|\mathcal{D}|}{2}}{\frac{1}{2} \left[ \sum_i \binom{\nu_i}{2} + \sum_j \binom{\tilde{\nu}_j}{2} \right] - \left[ \sum_i \binom{\nu_i}{2} \sum_j \binom{\tilde{\nu}_j}{2} \right] / \binom{|\mathcal{D}|}{2}}, \quad (15)$$

$$ACC = \max_{m \in \mathcal{M}} \frac{\sum_{i=1}^{|\mathcal{D}|} \mathbf{1}\{l_i == m(c_i)\}}{|\mathcal{D}|}. \quad (16)$$

Higher values evaluated by NMI, ARI, and ACC indicate better clustering quality. They are all equal to 1 when a perfect match is achieved in cluster assignments on a whole corpus. Both NMI and ARI penalize unnecessary splits of texts from the same true cluster to several inferred clusters. ARI further penalizes undesirable merging of texts from different true clusters into the same inferred cluster, making it a more rigorous metric.

### 5.4 Software Configuration

We conduct all the experiments on a server with 32 CPUs and 4 GPUs (GTX1080Ti). For all baselines, the software environments are configured according to their specific requirements. Extensively used in the literature, Scikit-learn is adopted for the implementations of K-means, HieClu, and LDA, the visualization tool t-SNE, and the metrics NMI, ARI, and ACC. Our models ARL and ARL-Adv are built based on Python and Tensorflow, with CUDA and cuDNN to enable the usage of GPUs. We release the Event dataset<sup>15</sup> described in Section 5.1. With the publication of this paper, the source code will be prepared for relevant studies.

### 5.5 Model Comparison (Q1)

We report the evaluation results in Table 3 for ARL, ARL-Adv, and all the baselines. The cluster numbers here are set to be the same as those in Table 1. From a whole perspective, all the methods do not perform the same w.r.t. the three metrics. For example, K-means(TF-IDF) gains better results than LDA w.r.t. NMI and ACC, but exhibits worse performance w.r.t. ARI. This

14. <https://radimrehurek.com/gensim/models/word2vec.html>

15. <https://github.com/CodeUpload-E/ARL-Adv>

TABLE 4: Ablation study of the proposed model.

Dataset Metrics	TREC			GoogleNews			Event			StackOverflow		
	NMI	ARI	ACC	NMI	ARI	ACC	NMI	ARI	ACC	NMI	ARI	ACC
ARL-Adv(no train w)	0.9001	0.7257	0.7631	0.8699	0.6483	0.6991	0.7039	0.3783	0.4450	0.3871	0.1819	0.2917
ARL-Adv(no train c)	0.9029	0.7755	0.7954	0.8783	0.7335	0.7688	0.7740	0.6230	0.6072	0.4787	0.2870	0.3859
ARL-Adv w/o $\mathcal{L}_1$	0.9279	0.8408	0.8368	0.9005	0.7881	0.8044	0.6826	0.3915	0.4631	0.3483	0.0942	0.2294
ARL-Adv w/o $\mathcal{L}_2$	0.9286	0.8441	0.8380	0.9033	0.8120	0.8240	0.8176	0.6244	0.6297	<b>0.4946</b>	<b>0.3071</b>	<b>0.4021</b>
ARL	0.9261	0.8428	0.8347	0.8995	0.7897	0.8057	0.8243	0.6078	0.6156	0.4796	0.2731	0.3781
ARL-Random	0.9250	0.8395	0.8330	0.8980	0.7813	0.7998	0.8174	0.6005	0.6064	0.4761	0.2564	0.3678
ARL-Adv(word)	0.9271	0.8458	0.8358	0.9014	0.7925	0.8074	0.8294	0.6387	0.6307	0.4793	0.2604	0.3684
ARL-Adv	<b>0.9305</b>	<b>0.8486</b>	<b>0.8402</b>	<b>0.9054</b>	<b>0.8214</b>	<b>0.8303</b>	<b>0.8450</b>	<b>0.6728</b>	<b>0.6607</b>	0.4871	0.3006	0.3943

phenomenon shows the necessity of adopting all the three metrics. Since they focus on different properties of the clustering results, using all of them can provide more comprehensive comparisons from distinct perspectives.

As expected, the models belonging to the conventional category (i) do not show competitive performance than other types of baselines in most cases, since they are not tailored for short texts and do not fully leverage the power of representation learning. Specifically, LDA works poorly here because assigning different clusters to words in the same short text can aggravate sparsity, as the corresponding results demonstrate. Although PCA converts the sparse representation of TF-IDF to continuous low-dimensional space, K-means(PCA) shows no improvements over K-means(TF-IDF). This result indicates that introducing representation that is independent of the model training process may not easily work. Besides, although the classical and simple model HieClu adopts another mechanism (i.e., agglomerative clustering), its performance is still at the same level as K-means.

Both DEC and VaDE are deep learning based approaches (ii), which can learn low-dimensional representations of targets using their original representations, such as the pixel values of images and TF-IDF based representations of texts. Therefore, though proposed for general clustering tasks, they may not be limited to texts. As shown in Table 3, they yield better results in most cases than the category of conventional methods. By further investigating the results of ClusterGAN, we find it does not exhibit the satisfied clustering performance. This might be attributed to the reason that the model is incapable of learning representations for discrete text tokens due to its GAN-based design.

Among the baselines originally proposed for short text clustering (iii), STCC performs not very well. This comparison shows that separating the representation learning and clustering into different stages tends to be sub-optimal. The reason might be that the representation learning process lacks proper guidance from the feedback of clustering if they are not learned together. Compared to STCC, GSDMM and BTM perform much better in most cases. In particular, GSDMM outperforms the general deep learning based clustering models in most cases. We attribute this phenomenon to the proper assumption in GSDMM that short texts are generated from only one topic, along with its proper approximation in the posterior distribution. STC exhibits good performance of TREC and StackOverflow (still worse than ARL-Adv) but not so good results on the other two datasets. It makes sense since STC cannot optimize word embeddings when training its clustering method.

To sum up, our full model ARL-Adv is superior since its improvements over the baselines are statistically significant on all the datasets, verified by t-test. This might be attributed to the following reasons: (1) ARL-Adv utilizes word embeddings

to obtain short text embeddings and learns word embeddings with the optimization of the whole model rather than learning it separately. (2) ARL-Adv combines the representation learning and short text clustering in an end-to-end learning fashion. (3) Robust adversarial training adopted in ARL-Adv can effectively improve clustering performance. The last part of Table 3 shows that ARL-Adv improves ARL consistently in the four datasets, especially for GoogleNews, Event, and StackOverflow, verifying the third reason. In the next section, we empirically demonstrate the benefit of learning word embeddings and cluster representations by the ablation study on the model.

## 5.6 Analysis of the Proposed Model (Q2)

Table 4 shows the results of the ablation study, so we can investigate the role of critical components in our proposed model. By first comparing ARL-Adv(no train w) and ARL-Adv(no train c) with ARL-Adv, we can see their performance drops significantly. This phenomenon shows the learning of word representations and cluster representations, accompanied by the automatic selection of clusters, is indeed indispensable. Besides, ARL-Adv(no train c) behaves better than ARL-Adv(no train w), indicating the training of word embeddings might be more critical for the model.

The second part of Table 4 shows the contributions of the pairwise-ranking loss and pointwise loss. For TREC, the two losses play similar roles, and removing either of them does not obviously damage the performance. On the contrary, the pairwise ranking loss is more crucial for the clustering performance on the other three datasets, and adding the pointwise loss could strengthen short text clustering in most cases — no improvement observed in StackOverflow. Actually, we could add a hyperparameter in Equation (8) and (9) to control the relative influence of the pointwise loss for a specific dataset.

The third part of Table 4 further explores the effect of adversarial cluster-level perturbations, by comparing it with cluster-level random perturbations (ARL-Random) and adversarial word-level perturbations (ARL-Adv(word)). Based on the results, we observe that: (1) ARL-Random obtains slightly worse results than ARL. This reveals that simply adding random perturbations without learning may not bring more useful information to the model. (2) ARL-Adv(word) performs marginally better than ARL in most cases, but not as well as ARL-Adv, especially for GoogleNews, Event, and StackOverflow. For example, compared to ARL-Adv(word), ARL-Adv gains the relative improvements of 2.8%, 4.8%, and 7.0% evaluated by ACC for the three datasets, respectively.

## 5.7 Qualitative Studies (Q3)

To explore the representations generated by ALR-Adv and other baselines, we use t-SNE [39] to visualize their short text embed-



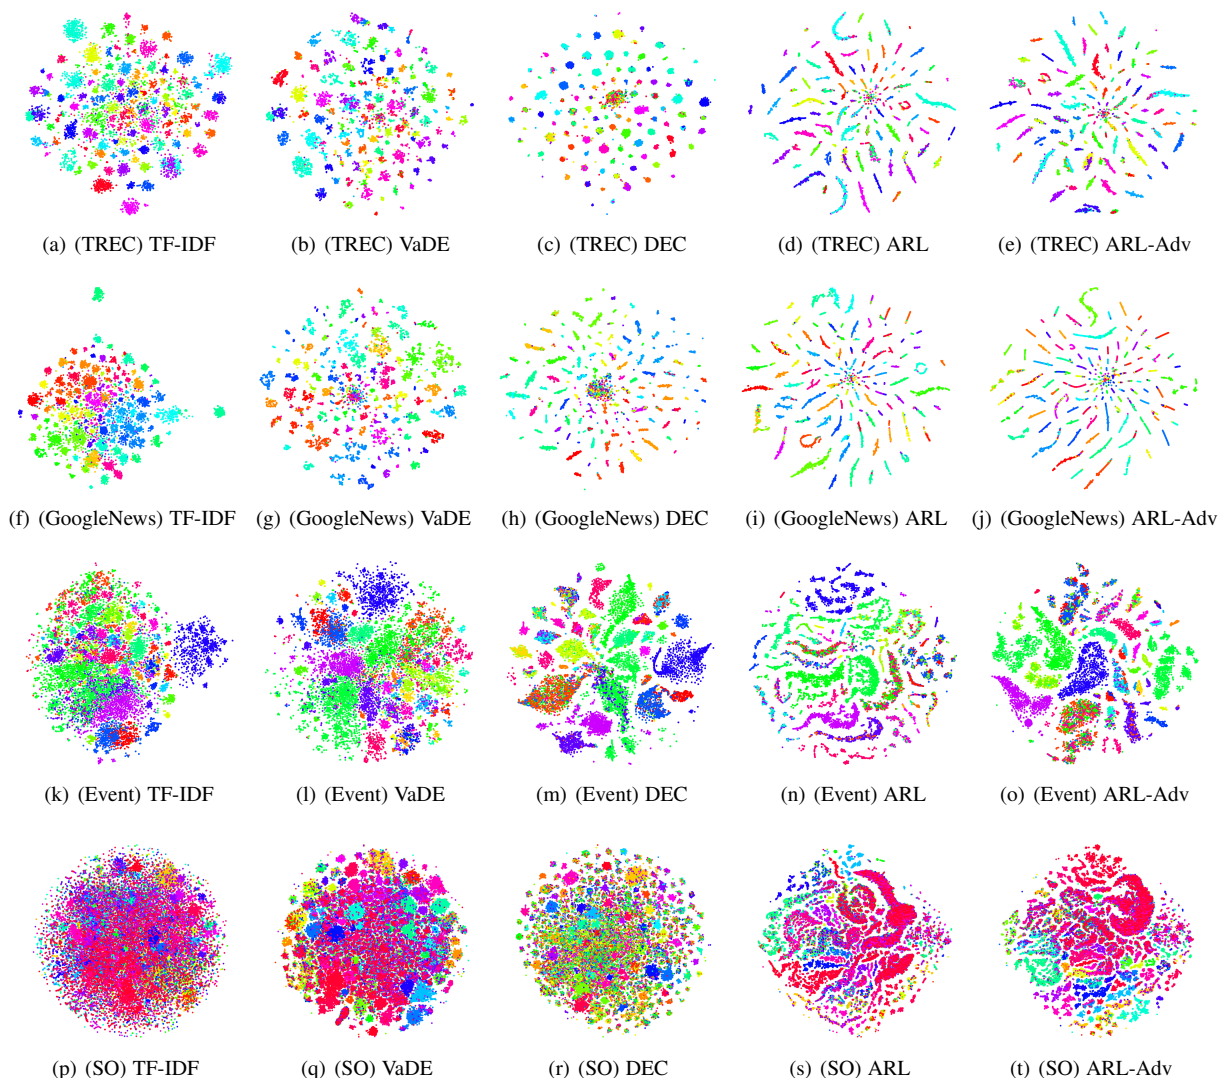


Fig. 2: Visualization of all short text representations in a 2-dimensional space using t-SNE. The quality of TF-IDF features, short text embeddings from VaDE, DEC, and the proposed models are investigated.

TABLE 5: Representative topical words of TREC.

Topic	ARL-Adv	GSDMM
1	card consolid credit debt loan unsecur advic settlement counsel creditor bailout	debt credit consolid card loan settlement relief bad negoti post unsecur combin lower
2	gifford gabriell buildup fluid congresswoman rehabilit icu recoveri brain hospit doctor	gifford gabriell rehab doctor recoveri brain news congresswoman fluid buildup
3	drone uav deliveri airspac southwest amazon chair requir commerci propos slice	drone amazon propos airspac deliveri commerci zone sky fli uav high-spe plan

TABLE 6: Representative topical words of Event.

Topic	ARL-Adv	GSDMM
1	daallo plane somalia land somal beachfront shabaab au hablod mogadishu airplan	somalia plane land emerg explos forc daallo airlin injur make passeng hole blast
2	plow plough bastilleday nice revel speed promenad driven franceattack truck	truck crowd attack prayformic innoc heart shatter tie enjoy parad plow leader
3	lahoreblast pakistanbomb lahor peshawar prayforlahor peshawarblast guess	lahoreblast islamabad pakistan pm armi punjab govt oper blast lahor minist call

dings in Figure 2, where one color denotes a cluster. We find that ARL-Adv and ARL generate highly separable semantic clusters, compared with high-dimensional yet sparse TF-IDF based features and low-dimensional embeddings from VaDE and DEC. This visualization explains the outstanding scores of ARL-Adv evaluated by the three metrics. DEC shows better clustering visualization than the other two baselines. However, the central areas in TREC and GoogleNews are somehow fuzzy, and clusters in Event and

StackOverflow are not separated so well. Even for humans, the clusters provided by ARL-Adv and ARL are easy to identify, exhibiting its potential as the basis or initialization for downstream applications that require insights into data manifold. Besides, there exist non-negligible differences between the visualizations of ARL-Adv and ARL, showing robust adversarial training indeed affects cluster representations. For example, embeddings provided by ARL-Adv tend to be more compact than those provided by




Fig. 3: Influence of  $\alpha$  on the model performance.




Fig. 4: Influence of  $\epsilon$  on the model performance.

ARL, which is more apparent for Event and StackOverflow.

We further evaluate the cluster-level quality by comparing the keywords obtained by ARL-Adv and GSDMM, presented in Table 5 and 6. One way of regarding a topic as being detected is to judge if it is held by most short texts belonging to a cluster. Hence, we can align clusters from different methods according to the detected topics. STCC, VaDE, and DEC may not afford such information since they do not directly model the correlations between words and clusters. While GSDMM chooses the words based on the learned word distributions for different topics, ARL-Adv selects representative words that have larger cosine similarity values with the learned cluster representations. The words that without identifiable meanings are discarded and the top listed words from the remaining are chosen.

Both ARL-Adv and GSDMM share lots of keywords in common, showing that they detect similar topics and reach consensus to a certain degree on words that better reflect the topics. Taking topic 2 of TREC as an example, words like *brain*, *fluid*, and *recoveri* imply the case of a cured illness, related to the entity *gifford* or *gabriell*. However, GSDMM only identifies *doctor*, while ARL-Adv further points out *hospit* and *icu*. The two words are even not ranked higher in the word lists of GSDMM, yet they are indeed highly correlated to the topic. For Event, we can observe a similar phenomenon that overlapped keywords still exist between the two methods, confirming their capability of providing an overview of a cluster. As revealed in topic 1 of Event, while

both of them succeed to identify *dallo*, *somalia*, and other words, ARL-Adv goes further and extracts highly correlated entities like *beachfront* and *shabaab*. These evidences consolidate our viewpoint that ARL-Adv can capture highly correlated keywords, which tend to be neglected by conventional topic models like GSDMM. Hence our model is welcomed by better semantic quality at the cluster level.

## 5.8 Hyper-parameter Analysis

### 5.8.1 Effect of $\alpha$

As described in Equation (11),  $\alpha$  adjusts the strength of  $J_2$  and thus constrains the contribution of adversarial perturbations to the final loss function. Figure 3 shows how the clustering performance of ARL-Adv alters along with the selection of  $\alpha$ , which helps us to understand the role of loss  $J_2$ . Note that  $\epsilon$  is set to the default value 50 as shown in Table 2. The figures show that a wide range of  $\alpha$ , from 0.2 to 8, can already provide a significant gain in most cases among the datasets and metrics. Moreover, selecting  $\alpha$  around 1.0 achieves better results than most values in TREC, GoogleNews, and StackOverflow. Thus reaching a balance between the loss function  $J_1(\mathbf{E}, \mathbf{C})$  and its adversarial version  $J_2(\mathbf{E}, \mathbf{C} + \Delta^C)$  is a preferable choice in ARL-Adv and we fix  $\alpha=1.0$  throughout the experiments. In summary, ARL-Adv relies, but not completely, on robust adversarial learning, showing that it appropriately and effectively takes the advantage of perturbations.




Fig. 5: Influence of the cluster number. The ratio of intended cluster number to the original topic number is used.

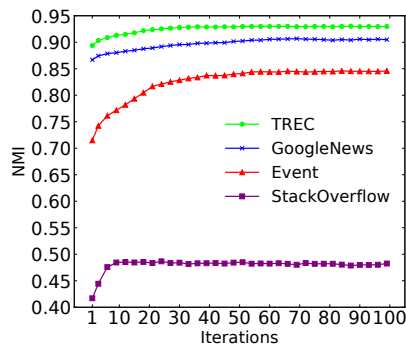


Fig. 6: Performance over iterations.

### 5.8.2 Effect of $\epsilon$

Figure 4 presents the performance of ARL-Adv when varying  $\epsilon$ . As defined in Equation (13),  $\epsilon$  constrains the norm of adversarial perturbations on cluster representations. The larger  $\epsilon$  is, the larger the adversarial perturbations will be during the training process. Under the setting of  $\alpha$  discussed above, we select some values for  $\epsilon$  in the range from 1 to 100, and show how it affects the clustering performance. We can infer from the figures that positive impact can be achieved locally when  $\epsilon$  falls in the range [10, 70] in most cases. Noticeably, while increasing  $\epsilon$  based on the chosen range indeed shows a growing tendency for Event, it does not apply to the other three datasets. Thus, we do not alter  $\epsilon$  to larger values since it may not reflect the universal property of ARL-Adv. Besides, varying  $\epsilon$  below 1 does not seem to be helpful, so we do not list the results here due to the limited space. Consequently, we can summarize that the restriction imposed on the norm of  $\Delta^C$  should lie in a proper scope to get reasonable results.

### 5.8.3 Effect of Cluster Number

Figure 5 shows how ARL-Adv adapts to the change of cluster number on the first three datasets w.r.t. NMI. The ratio of the predetermined cluster number to the original topic number is used for convenience. The initialization of cluster representations is the same as described in Section 5.2.2. We choose to compare ARL-Adv against GSDMM, which is one of the most competitive baseline, and K-means(PCA), which offers insights from the perspective of spectral clustering. For all the considered methods, a dramatic increase is found when the ratio goes from 0.2 to 0.4. This phenomenon is intuitive since the models are forced to merge texts from different topics into several clusters, which will almost inevitably lead to worse clustering results w.r.t. NMI. After

TABLE 7: Training time costs (time unit: s).

Data	ClusterGAN	DEC	VaDE	STC	ARL	ARL-Adv
TREC	3.7	1.1	2.2	1.1	1.7	2.1
SO	42.9	30.2	40.1	27.3	24.5	27.6

that, the evaluation scores continue to grow, and gradually get saturated. K-means(PCA) fails to keep its performance when the ratio gets larger. Therefore, spectral-based clustering algorithms might lack flexibility in the context of short text clustering, and turn out to heavily rely on human expertise in the datasets. On the other hand, while increasing the ratio does not help too much for both ARL-Adv and GSDMM, ARL-Adv outperforms GSDMM in all situations by a large margin. This might be attributed to that they tend to keep their assignment strategies still even when the cluster counts are more than sufficient. As such, ARL-Adv is not only effective, but also adaptive and robust enough to determine the proper mapping for texts when facing more options.

### 5.8.4 Effect of Iterations

We show the performance of the proposed ARL-Adv over iterations in Figure 6. The curves are averaged over 10 runs and we choose a few points at each interval for clarity. Due to the effect of adversarial perturbations, ARL-Adv experiences a relatively moderate growth in its performance. After 50 or more iterations, ARL-Adv starts to attain a stable state on all metrics, sharing a good convergence property.

## 5.9 Time Complexity Analysis

In this section, we move to compare the training time efficiencies for the adopted representation learning approaches on a single GPU. Probabilistic topic models like GSDMM are not included since their Gibbs sampling based optimization implementations do not leverage the computational power of GPU well. We also do not report the efficiencies of K-means algorithms since they demand multiple processors to speed up the solution of assigning each short text to its closest center. To make fair comparisons, we set the batch size to 64 uniformly when testing the time costs of the methods to be compared.

Table 7 shows the time costs needed for each training iteration in average, from which we have the following three key observations. First, the computational efficiencies of the proposed models are comparable with the existing approaches. Second, adopting robust adversarial training does not additionally incur a very heavy burden. Third, an approximate linear increase in computational

TABLE 8: Results of cluster-based retrieval. Noting that “—” denotes the original language model for retrieval without using clusters.

Clustering Method	Precision@50	
	Event	StackOverflow
—	0.6333	0.8561
K-means(PCA)	0.6357	0.8550
GSDMM	0.6423	0.8542
DEC	0.6426	0.8571
<b>ARL-Adv</b>	<b>0.6472</b>	<b>0.8647</b>

complexity of ARL and ARL-Adv is observed by the time cost and data size comparisons between TREC and StackOverflow (SO). Since both ARL and ARL-Adv require the same or less training iterations — usually less than 60 as shown in Figure 6 — for convergence compared to the other approaches through the experiments, we could conclude that their computational complexities are at a good level. The results meet expectation since the proposed models have simple architectures and the model parameters only involve word embeddings, cluster representations, and adversarial perturbations of clusters.

### 5.10 Application on Cluster-based Retrieval

To verify the effect of short-text clustering models on the downstream applications, we choose the cluster-based retrieval task [37] as a specific scenario for testing. The task assumes that if a document is relevant to a given query, then the document cluster it belongs to should be related to the query as well. Specifically, we select CBDM [37], a variant of language models with cluster smoothing, as the retrieval model to validate the contribution of text clusters to retrieval performance. The clusters we choose are produced by ARL-Adv and some representative baselines, including K-means(PCA), GSDMM, and DEC.

Table 8 presents the retrieval results on the Event and StackOverflow datasets evaluated by Precision@50, from which we observe that: (1) Using the clusters generated by our model ARL-Adv indeed boosts the performance of the downstream application, i.e., cluster-based retrieval. (2) The gains brought by ARL-Adv are the largest among the chosen clustering methods.

## 6 CONCLUSION

In this paper, we have developed the novel clustering model ARL-Adv. It fuses the short text representation learning and clustering in a unified model through the proposed cluster-level attention. Adversarial perturbations are further added to cluster representations, enhancing the robustness and effectiveness of model training through a minimax game. Extensive studies on the four real-life datasets show that ARL-Adv can achieve superior performance, even compared to the state-of-the-art methods for short text clustering and the recently developed deep learning based clustering models. Further analysis of ARL-Adv is performed to explain the contributions of its components.

## REFERENCES

[1] C. C. Aggarwal and C. Zhai. A survey of text clustering algorithms. In *Mining text data*, pages 77–128. Springer, 2012.

[2] M. Ailem, F. Role, and M. Nadif. Sparse poisson latent block model for document clustering. *IEEE TKDE*, 29(7):1563–1576, 2017.

[3] D. Bahdanau, K. Cho, and Y. Bengio. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.

[4] S. Banerjee, K. Ramanathan, and A. Gupta. Clustering short texts using wikipedia. In *SIGIR*, pages 787–788, 2007.

[5] Y. Bengio. Learning deep architectures for ai. *Foundations and trends® in Machine Learning*, 2(1):1–127, 2009.

[6] D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent dirichlet allocation. *JMLR*, 3(Jan):993–1022, 2003.

[7] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, et al. Language models are few-shot learners. *Arxiv*, 2020.

[8] C. Carpineto and G. Romano. Consensus clustering based on a new probabilistic rand index with application to subtopic retrieval. *IEEE TPAMI*, 34(12):2315–2326, 2012.

[9] H. Chen, M. Sun, C. Tu, Y. Lin, and Z. Liu. Neural sentiment classification with user and product attention. In *EMNLP*, pages 1650–1659, 2016.

[10] R. Das, M. Zaheer, and C. Dyer. Gaussian lda for topic models with word embeddings. In *ACL*, pages 795–804, 2015.

[11] J. Devlin, M. Chang, K. Lee, and K. Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. In *NAACL*, pages 4171–4186, 2019.

[12] W. Feng, C. Zhang, W. Zhang, J. Han, J. Wang, C. C. Aggarwal, and J. Huang. STREAMCUBE: hierarchical spatio-temporal hashtag clustering for event exploration over the twitter stream. In *ICDE*, pages 1561–1572, 2015.

[13] T. Fu, S. Tai, and H. Chen. Attentive and adversarial learning for video summarization. In *WACV*, pages 1579–1587, 2019.

[14] L. Gao, X. Li, J. Song, and H. T. Shen. Hierarchical lstms with adaptive attention for visual captioning. *IEEE TPAMI*, 42(5):1112–1131, 2020.

[15] K. Ghasedi, X. Wang, C. Deng, and H. Huang. Balanced self-paced learning for generative adversarial clustering network. In *CVPR*, pages 4391–4400, 2019.

[16] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *NIPS*, pages 2672–2680, 2014.

[17] I. J. Goodfellow, J. Shlens, and C. Szegedy. Explaining and harnessing adversarial examples. 2015.

[18] A. Hadifar, L. Sterckx, T. Demeester, and C. Develder. A self-training approach for short text clustering. In *Repl4NLP*, pages 1–6, 2019.

[19] W. Harchaoui, P.-A. Mattei, and C. Bouveyron. Deep adversarial gaussian mixture auto-encoder for clustering. In *ICLR Workshop*, 2017.

[20] R. He, W. S. Lee, H. T. Ng, and D. Dahlmeier. An unsupervised neural attention model for aspect extraction. In *ACL*, pages 388–397, 2017.

[21] X. He, Z. He, X. Du, and T. Chua. Adversarial personalized ranking for recommendation. In *SIGIR*, pages 355–364, 2018.

[22] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.

[23] R. Huang, G. Yu, Z. Wang, J. Zhang, and L. Shi. Dirichlet process mixture model for document clustering with feature partition. *IEEE TKDE*, 25(8):1748–1759, 2013.

[24] L. Hubert and P. Arabie. Comparing partitions. *Journal of classification*, 2(1):193–218, 1985.

[25] R. T. Ionescu and A. M. Butnaru. Vector of locally-aggregated word embeddings (VLAWE): A novel document-level representation. In *NAACL*, pages 363–369, 2019.

[26] A. K. Jain. Data clustering: 50 years beyond k-means. *Pattern Recognition Letters*, 31(8):651–666, 2010.

[27] Z. Jiang, Y. Zheng, H. Tan, B. Tang, and H. Zhou. Variational deep embedding: An unsupervised and generative approach to clustering. In *IJCAI*, pages 1965–1972, 2017.

[28] N. Kalchbrenner, E. Grefenstette, and P. Blunsom. A convolutional neural network for modelling sentences. In *ACL*, pages 655–665, 2014.

[29] Y. Kim. Convolutional neural networks for sentence classification. In *EMNLP*, pages 1746–1751, 2014.

[30] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.

[31] D. P. Kingma and M. Welling. Auto-encoding variational bayes. In *ICLR*, 2014.

[32] H. W. Kuhn. The hungarian method for the assignment problem. *Naval Research Logistics*, 2(1-2):83–97, 1955.

[33] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.

[34] Y. Li, C. Luo, and S. M. Chung. Text clustering with feature selection by using statistical data. *IEEE TKDE*, 20(5):641–652, 2008.



- [35] S. Liang, E. Yilmaz, and E. Kanoulas. Dynamic clustering of streaming short documents. In *SIGKDD*, pages 995–1004, 2016.
- [36] S. Liang, E. Yilmaz, and E. Kanoulas. Collaboratively tracking interests for user clustering in streams of short texts. *IEEE TKDE*, 31(2):257–272, 2019.
- [37] X. Liu and W. B. Croft. Cluster-based retrieval using language models. In *SIGIR*, pages 186–193, 2004.
- [38] Q. Ma, L. Yu, S. Tian, E. Chen, and W. W. Y. Ng. Global-local mutual attention model for text classification. *IEEE/ACM TASLP*, 27(12):2127–2139, 2019.
- [39] L. v. d. Maaten and G. Hinton. Visualizing data using t-sne. *JMLR*, 9(Nov):2579–2605, 2008.
- [40] C. D. Manning, H. Schütze, and P. Raghavan. *Introduction to information retrieval*. Cambridge university press, 2008.
- [41] M. Mathioudakis and N. Koudas. Twittermonitor: trend detection over the twitter stream. In *SIGMOD*, pages 1155–1158, 2010.
- [42] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean. Distributed representations of words and phrases and their compositionality. In *NIPS*, pages 3111–3119, 2013.
- [43] T. Miyato, A. M. Dai, and I. Goodfellow. Adversarial training methods for semi-supervised text classification. In *ICLR*, 2017.
- [44] T. Miyato, S.-i. Maeda, M. Koyama, K. Nakae, and S. Ishii. Distributional smoothing with virtual adversarial training. In *ICLR*, 2016.
- [45] N. Mrabah, M. Bouguessa, and R. Ksantini. Adversarial deep embedded clustering: on a better trade-off between feature randomness and feature drift. *Arxiv*, 2019.
- [46] S. Mukherjee, H. Asnani, E. Lin, and S. Kannan. Clustergan: Latent space clustering in generative adversarial networks. In *AAAI*, pages 4610–4617, 2019.
- [47] V. Premachandran and A. L. Yuille. Unsupervised learning using generative adversarial training and clustering. *Arxiv*, 2016.
- [48] A. Rangrej, S. Kulkarni, and A. V. Tendulkar. Comparative study of clustering techniques for short text documents. In *WWW*, pages 111–112, 2011.
- [49] A. Ritter, E. Wright, W. Casey, and T. Mitchell. Weakly supervised extraction of computer security events from twitter. In *WWW*, pages 896–905, 2015.
- [50] M. Rosen-Zvi, T. L. Griffiths, M. Steyvers, and P. Smyth. The author-topic model for authors and documents. In *UAI*, pages 487–494, 2004.
- [51] R. Socher, A. Karpathy, Q. V. Le, C. D. Manning, and A. Y. Ng. Grounded compositional semantics for finding and describing images with sentences. *TACL*, 2(1):207–218, 2014.
- [52] A. Strehl and J. Ghosh. Cluster ensembles—a knowledge reuse framework for combining multiple partitions. *JMLR*, 3(Dec):583–617, 2002.
- [53] D. Tang, B. Qin, and T. Liu. Document modeling with gated recurrent neural network for sentiment classification. In *EMNLP*, pages 1422–1432, 2015.
- [54] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need. In *NIPS*, pages 5998–6008, 2017.
- [55] W. Wang, W. Zhang, J. Wang, J. Yan, and H. Zha. Learning sequential correlation for user generated textual content popularity prediction. In *IJCAI*, pages 1625–1631, 2018.
- [56] W. Wei, B. Xi, and M. Kantarcioglu. Adversarial clustering: A grid based clustering algorithm against active adversaries. *Arxiv*, 2018.
- [57] S. Xiao, J. Yan, M. Farajtabar, L. Song, X. Yang, and H. Zha. Learning time series associated event sequences with recurrent point process networks. *IEEE TNNLS*, 30(10):3124–3136, 2019.
- [58] J. Xie, R. Girshick, and A. Farhadi. Unsupervised deep embedding for clustering analysis. In *ICML*, pages 478–487, 2016.
- [59] Q. Xipeng, S. TianXiang, X. Yige, S. Yunfan, D. Ning, and H. Xuanjing. Pre-trained models for natural language processing: A survey. *SCIENCE CHINA Technological Sciences*, 2020.
- [60] J. Xu, P. Wang, G. Tian, B. Xu, J. Zhao, F. Wang, and H. Hao. Short text clustering via convolutional neural networks. In *NAACL*, pages 62–69, 2015.
- [61] J. Xu, B. Xu, P. Wang, S. Zheng, G. Tian, and J. Zhao. Self-taught convolutional neural networks for short text clustering. *Neural Networks*, 88:22–31, 2017.
- [62] X. Yan, J. Guo, Y. Lan, and X. Cheng. A bitern topic model for short texts. In *WWW*, pages 1445–1456, 2013.
- [63] Z. Yan, Y. Guo, and C. Zhang. Deep defense: Training dnns with improved adversarial robustness. In *NeurIPS*, pages 417–426, 2018.
- [64] Z. Yang, D. Yang, C. Dyer, X. He, A. J. Smola, and E. H. Hovy. Hierarchical attention networks for document classification. In *NAACL*, pages 1480–1489, 2016.
- [65] S. Yeung, O. Russakovsky, N. Jin, M. Andriluka, G. Mori, and L. Fei-Fei. Every moment counts: Dense detailed labeling of actions in complex videos. *IJCV*, 126(2-4):375–389, 2018.
- [66] J. Yin, D. Chao, Z. Liu, W. Zhang, X. Yu, and J. Wang. Model-based clustering of short text streams. In *SIGKDD*, pages 2634–2642, 2018.
- [67] J. Yin and J. Wang. A dirichlet multinomial mixture model-based approach for short text clustering. In *SIGKDD*, pages 233–242, 2014.
- [68] L. Yu, W. Zhang, J. Wang, and Y. Yu. Seqgan: Sequence generative adversarial nets with policy gradient. In *AAAI*, pages 2852–2858, 2017.
- [69] Y. Yu and W. Zhou. Mixture of gans for clustering. In *IJCAI*, pages 3047–3053, 2018.
- [70] H. Zhang, I. J. Goodfellow, D. N. Metaxas, and A. Odena. Self-attention generative adversarial networks. In *ICML*, pages 7354–7363, 2019.



**Wei Zhang** received his Ph.D. degree in computer science and technology from Tsinghua university, Beijing, China, in 2016. He is currently an associate researcher in the School of Computer Science and Technology, East China Normal University, Shanghai, China. His research interests mainly include data mining and machine learning applications. He has published more than 40 journal and conference papers in these areas.



**Chao Dong** is currently pursuing the master degree at East China Normal University, Shanghai, China. His main research interests include text mining and recommendation.



**Jianhua Yin** is currently working as an assistant professor in the Department of Computer Science and technology, Shandong University, Qingdao, China. He received his Ph.D. degree in Computer Science from Tsinghua University, Beijing, China in 2017. His main research interests include text mining, community question answering, and recommendation.



**Jianyong Wang** received the PhD degree in computer science from the Institute of Computing Technology, Chinese Academy of Sciences, in 1999. He is currently a professor in the Department of Computer Science and Technology, Tsinghua University, Beijing, China, and also with the Jiangsu Collaborative Innovation Center for Language Ability, Jiangsu Normal University, Xuzhou, China. His research interests mainly include data mining and Web information management. He has co-authored more than 60 papers in some leading international conferences and some top international journals. He is serving or has served as a PC member for some leading international conferences, such as SIGKDD, VLDB, ICDE, WWW, and an associate editor of the IEEE Transactions on Knowledge and Data Engineering and the ACM Transactions on Knowledge Discovery from Data. He is a fellow of the IEEE.