

EPiDA: An Easy Plug-in Data Augmentation Framework for High Performance Text Classification

Minyi Zhao¹ Lu Zhang¹ Yi Xu¹

Jiandong Ding² Jihong Guan³ Shuigeng Zhou^{1*}

¹Shanghai Key Lab of Intelligent Information Processing, and School of

Computer Science, Fudan University, China

²Alibaba Group ³Tongji University, China

^{1,2}{zhaomy20, l_zhang19, yxu17, jdding, sgzhou}@fudan.edu.cn ³jhguan@tongji.edu.cn

Abstract

Recent works have empirically shown the effectiveness of data augmentation (DA) for NLP tasks, especially for those suffering from data scarcity. Intuitively, given the size of generated data, their *diversity* and *quality* are crucial to the performance of targeted tasks. However, to the best of our knowledge, most existing methods consider only either the *diversity* or the *quality* of augmented data, thus cannot fully tap the potential of DA for NLP. In this paper, we present an easy and plug-in data augmentation framework EPiDA to support effective text classification. EPiDA employs two mechanisms: *relative entropy maximization* (REM) and *conditional entropy minimization* (CEM) to control data generation, where REM is designed to enhance the diversity of augmented data while CEM is exploited to ensure their semantic consistency. EPiDA can support efficient and continuous data generation for effective classifier training. Extensive experiments show that EPiDA outperforms existing SOTA methods in most cases, though not using any agent network or pre-trained generation network, and it works well with various DA algorithms and classification models. Code is available at <https://github.com/zhaominyiz/EPiDA>.

1 Introduction

Data augmentation (DA) is widely-used in classification tasks (Shorten and Khoshgoftaar, 2019; Feng et al., 2021; Zhang et al., 2021). In computer vision (CV), (Krizhevsky et al., 2012; Chatfield et al., 2014; Szegedy et al., 2015) adopt strategies like flipping, cropping, tilting to perform DA. In natural language processing (NLP), (Xie et al., 2017; Coulombe, 2018; Niu and Bansal, 2018; Wei and Zou, 2019) find that native augmentation skills such as spelling errors, synonym replacement, deleting and swapping, can bring considerable performance improvement. All these methods use

various transformations for data augmentation, but they do not achieve equal success in different NLP tasks (Yang et al., 2020). Sometimes, they fail to guarantee semantic consistency, and may even bring semantic errors that are harmful to classification. The reason lies in that data augmentation for NLP is in discrete space, so it can easily incur large deviation of semantics (*e.g.* in sentiment classification task, deleting emotional words from a sentence will make its meaning completely different).

Generally, given the size of generated data, their *diversity* and *quality* are crucial to the performance of targeted tasks (Ash et al., 2019). Recent works have begun to emphasize the diversity or quality of augmented data. For example, in CV, AA (Cubuk et al., 2019), Fast-AA (Lim et al., 2019) and LTA (Luo et al., 2020) employ agent networks to learn how to enhance diversity. In NLP, language models are widely used to control generation quality, including Back-translation (Sennrich et al., 2016; Yu et al., 2018), Seq2seq models (Kobayashi, 2018; Kumar et al., 2019; Yang et al., 2020), GPT-2 (Radford et al., 2019; Anaby-Tavor et al., 2020; Quteineh et al., 2020; Liu et al., 2020) and T5 (Dong et al., 2021). In addition, some works (Morris et al., 2020) in NLP utilize adversarial augmentation to enrich the diversity of the samples. However, to the best of our knowledge, most existing works consider only either the quality or the diversity of augmented data, so cannot fully exploit the potential of data augmentation for NLP tasks. Besides, recent existing DA methods for NLP tasks usually resort to pre-trained language models, are extremely inefficient due to huge model complexity and tedious finetuning, which limits the scope of their applications.

In this paper, we propose a new data augmentation framework for text classification. This framework is called EPiDA (the abbreviation of Easy Plug-in Data Augmentation), which employs two mechanisms to control the diversity and quality

*Corresponding author.




Figure 1: The pipeline of EPiDA. EPiDA first evaluates the candidate samples from the perspectives of diversity and quality, then takes the samples of high diversity and quality as training samples. It can work with various DA algorithms and classification models.

of augmented data: *relative entropy maximization* (REM) and *conditional entropy minimization* (CEM), where the former is for boosting diversity while the latter for ensuring quality. Fig. 1 shows the pipeline of EPiDA. EPiDA consists of a DA algorithm, a classifier, and a *Sample Evaluation And Selection* (SEAS) module. SEAS works with the DA algorithm and the classifier, and evaluates the candidate samples with the feedback of the classifier. With REM and CEM, SEAS can select samples of high diversity and quality to train the classifier continuously until the model converges.

The main contributions of this paper are as follows:

1. We propose an easy plug-in data augmentation framework EPiDA for text classification. EPiDA can work with various existing DA algorithms and classification models, it is general, efficient, and easy-to-deploy.
2. We design two mechanisms *relative entropy maximization* (REM) and *conditional entropy minimization* (CEM) to boost the diversity and quality of augmented data simultaneously in an explicit and controllable way.
3. We conduct extensive experiments to evaluate EPiDA. Experimental results show that EPiDA outperforms existing DA methods, and works well with different DA algorithms and classification models.

The rest of this paper is organized as follows: Sec. 2 reviews related work and highlights the differences between our work and major existing methods. Sec. 3 introduce our method in details. Sec. 4 presents the results of performance evaluation, and Sec. 5 concludes the paper.

2 Related Work

In this section, we first review the related work of DA for NLP, then expound the differences between our method and the major existing ones. According to the methodology of data generation, existing methods can be categorized into three types: rule-based, interpolation-based, and model-based, respectively.

2.1 Rule-Based Methods

These works use easy and predetermined transformations without model components. (Kolomiyets et al., 2011; Zhang et al., 2015; Wang and Yang, 2015) use synonyms to replace words. EDA (Wei and Zou, 2019) and AEDA (Karimi et al., 2021) introduce random insertions, swaps, and deletions. Xie et al. (2017) employed spelling errors to augment sentences. Şahin and Steedman (2018) conducted sentence rotating via dependency tree morphing. Wei et al. (2021) proposed a multi-task view of DA. SUB² (Shi et al., 2021) generates new examples by substituting substructures via constituency parse trees. Although these methods are easy to implement, they do not consider controlling data quality and diversity.

2.2 Interpolation-Based Methods

MIXUP (Zhang et al., 2017) pioneers this type of works by interpolating the input and labels of two or more real examples. Recently, many MIXUP strategies (Verma et al., 2019; Yun et al., 2019) were proposed in CV. Due to the discrete nature of inputs of NLP tasks, such methods can be applied to NLP tasks only via padding and mixing embeddings or higher hidden layers (Chen et al., 2020; Si et al., 2021).

2.3 Model-Based Methods

Seq2seq and language models have been used to generate high quality samples. Among these approaches, Back-translation (Sennrich et al., 2016; Yu et al., 2018) translates sentences into another language and then translates it back to the original language. RNNs and transformers are used to reconstruct sub-parts of real data with contextual information (Kobayashi, 2018; Gao et al., 2019; Yang et al., 2020). Recently, we have witnessed the great success of large-scale pre-trained language models (PLMs) such as BERT (Devlin et al., 2019), XLNet (Yang et al., 2019), GPT-2 (Radford et al., 2019) in NLP tasks. These state-of-the-art PLMs

Method	Div	Qua	LM	FB
AA (Cubuk et al., 2019)	✓	×	✓	✓
EDA (Wei and Zou, 2019)	✓	×	×	×
DataBoost (Liu et al., 2020)	×	✓	✓	×
LearnDA (Zuo et al., 2021)	✓	✓	✓	×
VDA (Zhou et al., 2021)	✓	✓	✓	×
Ours EPiDA	✓	✓	×	✓

Table 1: A qualitative comparison between EPiDA and major existing methods from four aspects: whether controlling the diversity and quality of the augmented data, whether using language model or agent network and whether using the feedback of the classifier.

are also widely used to augment samples (Ng et al., 2020; Nie et al., 2020; Anaby-Tavor et al., 2020; Quteineh et al., 2020; Liu et al., 2020; Dong et al., 2021). For example, DataBoost (Liu et al., 2020) develops a reinforcement learning strategy to guide the conditional generation without changing the architecture of GPT-2. Besides, adversarial augmentation (*i.e.*, attack, GANs) are also used to enrich the diversity of the generated samples (Morris et al., 2020; Simoncini and Spanakis, 2021). Although model-based methods can control generation quality well via PLMs, they are computationally inefficient, which limits their applications.

2.4 Differences between EPiDA and Existing Methods

To expound the differences between EPiDA and typical existing methods, in Tab. 1 we present a qualitative comparison from four dimensions: whether controlling the diversity and quality of the augmented data, whether using pre-trained model (language model or agent network), and whether using the feedback from the classifier.

As shown in Tab. 1, among the existing methods, most control only either diversity (*e.g.* AA and EDA) or quality (*e.g.* DataBoost) of augmented data, thus cannot completely leverage the potential of data augmentation. And most use language model or agent network, which is beneficial to data quality but also inefficient. Only the recent LearnDA (Zuo et al., 2021) and VDA (Zhou et al., 2021) consider both diversity and quality, and only AA uses feedback of the classifier. Our EPiDA addresses both diversity and quality of augmented data via the feedback of the classifier in an explicit and controllable way, without the help of any additional model components, which makes it not only more effective but also more efficient.

Note that in addition to the differences listed in Tab. 1, our method EPiDA differs from LearnDA in at least three other aspects: 1) LearnDA employs perplexity score (PPL) and cosine similarity to measure diversity and quality respectively, while EPiDA adopts two mechanisms *relative entropy maximization* (REM) and *conditional entropy minimization* (CEM) to control diversity and quality, which is theoretically more rational and solid. 2) LearnDA is for event causality identification, while EPiDA is mainly for text classification. 3) LearnDA needs knowledge guidance, while EPiDA does not. These make it difficult to evaluate LearnDA in our experimental settings. Thus, we do not conduct performance comparison between EPiDA and LearnDA. Nevertheless, in our ablation study, we replace REM and CEM with PPL and cosine similarity in EPiDA, and our experimental results show that EPiDA with REM and CEM performs better than that with PPL and cosine similarity. Besides, comparing with VDA that requires PLM to provide substitution probability, EPiDA is free of PLMs, and is more effective, efficient and practical.

3 Method

As shown in Fig. 1, EPiDA consists of three components: a DA algorithm T , a classifier or classification model C , and a *Sample Evaluation and Selection* (SEAS) module that is the core component of EPiDA. Generally, the DA algorithm and the classifier can be any of existing DA algorithms and classifiers. With the feedback of the classifier, SEAS evaluates candidate samples generated by the DA algorithm in terms of *diversity* and *quality* via the *Relative Entropy Maximization* (REM) mechanism and the *Conditional Entropy Minimization* (CEM) mechanism, and outputs the qualified samples to further train the classifier. So EPiDA can serve as a plug-in component to boost existing DA algorithms for training better target models.

3.1 The Rationale to Control DA

Consider a classification task with a dataset X of n samples: $X = \{(x_1, y_1), (x_2, y_2), \dots, (x_n, y_n)\}$. Here, x_i is a sample, y_i is its label. The loss function is

$$L_o(\omega) = \frac{1}{n} \sum_{i=1}^n l(\omega^\top \phi(x_i); y_i). \quad (1)$$

where $\phi: \mathbb{R}^d \rightarrow \mathbb{R}^D$ is a finite-dimensional feature map, $\omega \in \mathbb{R}^D$ means learnable parameters, and l

can be a common loss function like cross-entropy.

Now we employ a DA algorithm T to conduct augmentation for each sample in X . Let t_i^j be the j -th sample generated by T with x_i as input, and m samples are generated from x_i , the loss function for the generated samples can be written as

$$L_g(\omega) = \frac{1}{n} \sum_{i=1}^n \frac{1}{m} \sum_{j=1}^m l(\omega^\top \phi(t_i^j); y_i). \quad (2)$$

Here, we assume 1) t_i^j and x_i have the same label y_i , so we can use y_i to optimize the new loss function; 2) Data augmentation does not significantly change the feature map ϕ , that is, augmentation can maintain semantic consistency of the sample space. Now we combine the augmented samples into the original samples, thus the total loss function of EPiDA can be written as follows:

$$L(\omega) = L_o(\omega) + L_g(\omega). \quad (3)$$

Recall that we use the feedback of the classifier C to select samples. Specifically, we use the original training samples X to pre-train the classifier C , and for each generated sample t_i^j , the feedback signal about t_i^j from the classifier is used for evaluating t_i^j . When the generation process is over, all generated samples $\{t_i^j\}$ are used to train C again.

First, we consider how to generate samples of high diversity. Intuitively, generated samples should be different from the original samples. Recalling that the classifier C is pretrained by X , so for generated sample t_i^j , its loss $l(\omega^\top \phi(t_i^j); y_i)$ should be large. In this sense, given the classifier C (ω is fixed), we select samples that meet the following objective function:

$$\max_{t_i^j} L_g(\omega, \phi(t_i^j)), \quad (4)$$

which means that we are to generate ‘‘hard’’ samples for the classifier to cope with.

Second, we consider how to control the quality of augmented data. Recall that we assume for each augmented sample t_i^j , its label y_i keeps unchanged, so we can use the original label to evaluate the loss function. However, due to the discrete nature of language, it is nontrivial for augmented samples to meet this assumption. Taking the sentiment analysis task for example, suppose we use EDA (Wei and Zou, 2019) to augment x_i =‘‘you’ll probably love it’’, EDA may delete the word ‘‘love’’. Obviously,

the resulting sentence breaks the semantic consistency. To guarantee semantic consistency, we limit the semantic deviation of $\phi(t_i^j)$ from $\phi(x_i)$. Let M and ρ be a metric function to measure semantic difference between samples and a threshold respectively, we impose the following constraint on $\phi(t_i^j)$:

$$|M(\omega^\top \phi(t_i^j), \omega^\top \phi(x_i))| \leq \rho. \quad (5)$$

Thus, we can enhance data diversity by optimizing Eq. (4), and improve data quality using Eq. (5). The problem turns to solve Eq. (4) and Eq. (5).

3.2 Relative Entropy Maximization

We rewrite the objective function in Eq. (4) via:

$$\begin{aligned} L_g(\omega, \phi(t_i^j)) &= \frac{1}{nm} \sum_{i=1}^n \sum_{j=1}^m l(\omega^\top \phi(t_i^j); y_i) \\ &= \frac{1}{nm} \sum_{i=1}^n \sum_{j=1}^m D(p(\omega^\top \phi(t_i^j)), p(y_i)) + H(p(y_i)), \end{aligned} \quad (6)$$

where p , H , D indicate probability distribution, *Shannon entropy*, and *relative entropy* respectively, and actually $H(p(y_i)) = 0$ since $p(y_i)$ is a one-hot vector. According to Eq. (6), we try to augment samples with large relative entropy under the given labels. Thus, we call this method *relative entropy maximization* (REM) mechanism. As relative entropy measures the difference between the two distributions $p(\omega^\top \phi(t_i^j))$ and $p(y_i)$, the larger the difference is, the more diverse the augmented sample is. Therefore, we define the *diversity score* s_{div}^{ij} of augmented sample t_i^j as follows:

$$s_{div}^{ij} = D(p(\omega^\top \phi(t_i^j)), p(y_i)). \quad (7)$$

3.3 Conditional Entropy Minimization

We use *conditional entropy* as the metric function in Eq. (5) to constrain the semantic deviation of $\phi(t_i^j)$ from $\phi(x_i)$, i.e., $M(\cdot, \cdot) := H(\cdot|\cdot)$. Then, Eq. (5) can be rewritten to

$$H(p(\omega^\top \phi(t_i^j))|p(\omega^\top \phi(x_i))) \leq \rho. \quad (8)$$

where $H(\cdot, \cdot)$ is *conditional entropy*. Furthermore, to meet Eq. (8), we select samples $\{t_i^j\}$ by solving the following optimization problem:

$$\min_{t_i^j} H(p(\omega^\top \phi(t_i^j))|p(\omega^\top \phi(x_i))). \quad (9)$$

We call this *conditional entropy minimization* (CEM) mechanism. The smallest value of

$H(p(\omega^\top \phi(t_i^j))|p(\omega^\top \phi(x_i)))$ is 0, indicating that given $p(\omega^\top \phi(x_i))$, $p(\omega^\top \phi(t_i^j))$ is exactly predictable. Eq. (9) can also be expanded to the difference between Shannon entropy H and mutual information I , i.e., $H(X|Y)=H(X)-I(X, Y)$. In other words, CEM minimizes the entropy of the selected sample t_i^j and maximizes the mutual information between t_i^j and the original sample x_i , which means that CEM tries to augment samples of high prediction probability and high similarity with the original sample. As in REM, we define the *quality score* s_{qua}^{ij} of augmented sample t_i^j as

$$s_{qua}^{ij} = -H(p(\omega^\top \phi(t_i^j))|p(\omega^\top \phi(x_i))). \quad (10)$$

Algorithm 1: EPiDA Data Augmentation.

Input: Classification model ω , input sample x_i and its label y_i , DA algorithm T , augmentation number m and amplification factor K .

Output: m augmented samples.

// assign the set $T(x_i)$ of mK candidates to array \hat{t}_i

$\hat{t}_i \leftarrow T(x_i)$;

$s_{div}, s_{qua}, s_{tot} = \mathbb{R}^{K*m}, \mathbb{R}^{K*m}, \mathbb{R}^{K*m}$;

for $j = 1, 2, \dots, K * m$ **do**

 Calculate s_{div}^{ij} via Eq. (7) ;

 Calculate s_{qua}^{ij} via Eq. (10) ;

end

Take Min_Max_Norm for s_{div} and s_{qua} ;

$s_{tot} = s_{div} + s_{qua}$;

// find the subscripts of the top m small elements

$id = \text{argtopm}(-s_{tot})$;

Return $\hat{t}_i[id]$;

3.4 Algorithm and Implementation

The procedure of EPiDA is presented in Alg. 1. For each input sample x_i , EPiDA outputs m augmented samples. First, we employ T to generate $K * m$ candidate augmented samples for x_i , where K is a hyperparameter to amplify the number of candidate samples, which is called *amplification factor*. Then, for each augmented sample, we use REM and CEM to evaluate its diversity score (s_{div}) and quality score (s_{qua}), respectively. Next, we adopt Min_Max_Norm to make s_{div} and s_{qua} fall in $[0,1]$. After that, we add them together as the overall score of the sample, and sort all the augmented samples in descending order according to

Sentence	s_{div}	s_{qua}	s_{tot}
Go Set a Watchman comes out Tuesday and I'm really <u>excited</u> for it	0.00	1.00	1.00
Go Set a Watchman comes out Tuesday and I'm really <u>mad</u> for it	0.96	0.03	0.99
Go Set a Watchman out Tuesday and I'm really <u>excited</u> for it	0.05	0.92	0.97
Go Set a security guard comes out Tuesday and I'm really <u>excited</u> for it	0.86	0.15	1.01

Table 2: An example from the sentiment analysis task to demonstrate the effect of REM and CEM. The first row is the original sample, the rest are the augmented samples. Underlined are salient words.

their scores. Finally, we take the top m samples from all the $K * m$ candidate samples as the output, and utilize them to train the classifier.

By nature, the goals of REM and CEM are conflicting, i.e., a sample of high diversity is more probably of low quality, and vice versa. We give an example in Tab. 2 to demonstrate this point. REM encourages to change salient words, which is prone to break the semantic consistency (see the 3rd row, “excited” is changed to “mad”, leading to large diversity score but small quality score). However, CEM tends to make the augmented samples keep semantic consistency, i.e., has large quality score but small diversity score (see the 4th row, “comes” is deleted). By jointly considering REM and CEM, satisfactory samples with balanced diversity and quality can be found (see the 5th row).

Besides, the calculation of s_{div} and s_{qua} requires the feedback of the classifier, so we first pre-train the classifier using the original samples, then with EPiDA we can generate samples of high diversity and quality for the classifier continuously.

4 Performance Evaluation

In this section, we conduct extensive experiments to evaluate EPiDA, including performance comparison with SOTA methods, performance evaluation when working with different DA algorithms and classification models, ablation study, and qualitative visualization of samples augmented by EPiDA.

4.1 Datasets and Settings

Datasets for five different tasks are used in our experiments: *Question Classification* (Li and Roth, 2002) (TREC, $N=5,452$), *News Classification* (Zhang et al., 2015) (AGNews, $N=120,000$), *Tweets Sentiment Analysis* (Rosenthal et al., 2017)

Method	Sentiment			Irony			Offense		
	10%	40%	PPL	10%	40%	PPL	10%	40%	PPL
EDA (randomly delete, swap etc.)	0.560	0.608	41.22	0.530	0.515	76.07	0.637	0.629	37.37
Contextual Word Embs Aug. (insert, replace using Bi-RNN LM)	0.610	0.627	1043.18	0.518	0.593	1146.40	0.663	0.713	1729.62
Back-translation Aug. (Eng. → Fr. → Eng. as aug. text)	0.617	0.620	474.29	0.520	0.541	423.32	0.655	0.724	345.23
Data Boost (RL-guided conditional generation)	0.591	0.642	56.23	0.591	0.639	77.40	0.695	0.784	35.18
Ours EPiDA: REM only (EDA as the DA algorithm)	0.619	0.650	66.86	0.624	0.665	77.09	0.662	0.673	81.13
Ours EPiDA: CEM only (EDA as the DA algorithm)	0.629	0.659	8.10	0.629	0.666	12.11	0.668	0.670	8.57
Ours EPiDA: REM+CEM (EDA as the DA algorithm)	0.639	0.659	25.40	0.651	0.687	53.17	0.680	0.687	32.56

Table 3: Performance comparison with existing augmentation methods. 10%: 10% original data + 30% augmented data ($m = 3$); 40%: 40% original data + 40% augmented data ($m = 1$). We report the F1 score of the BERT classifier averaged over five repeated experiments on each dataset. We also report the perplexity score (PPL) of 10,000 randomly sampled data augmented by each method, where PPL is evaluated by the kenLM language model trained on the training data of each task.

(**Sentiment**, $N=20,631$), *Tweets Irony Classification* (Van Hee et al., 2018) (**Irony**, $N=3,817$) and *Tweets Offense Detection* (Founta et al., 2018) (**Offense**, $N=99,603$), where N is the number of training samples. To fully demonstrate the performance of data augmentation, we use only part of each dataset. In the following experiments, the percentage (%) that follows the task name means the ratio of training data used from each dataset, e.g. Irony 1% means that 1% of the dataset is used. Macro-F1 (F1 for binary tasks) is used as performance metric, and all the experiments are repeated five times. The amplification factor K is set to 3.

4.2 Comparing with SOTA Methods

Here we carry out performance comparison with major SOTA methods to show the superiority of EPiDA on three datasets: **Sentiment**, **Irony** and **Offense**. For a fair comparison, we strictly follow the experimental setting of DataBoost (Liu et al., 2020): we do only one round of augmentation to ensure that the number of samples of our method is consistent with that of the other methods, and use BERT as the classifier. We use the widely used EDA as the DA algorithm of EPiDA. We do not use DataBoost because it is not yet open-sourced. The experimental results are presented in Tab. 3.

From Tab. 3, we can see that 1) with the help of EPiDA, the performance of EDA is greatly improved. In particular, comparing with the original EDA, EPiDA gets performance improvement of 14.1%, 8.39%, 22.83%, 33.40%, 6.75% and 9.22% in six task settings, respectively. 2) Our method outperforms DataBoost in four settings. In particular, EPiDA+EDA gets performance improvement of 8.12%, 2.65%, 10.15% and 6.99% in various settings of the **Sentiment** and **Irony** tasks. 3) The variants of EPiDA that utilize only REM or CEM to enhance diversity or quality are inferior to using both, which demonstrates the effectiveness of joint enhancement. 4) DataBoost performs better in the **Offense** task, the reason lies in that DataBoost can create novel sentences from **Offense** (a relatively huge corpus) via GPT-2, while EDA only conducts word-level augmentation, which limits EPiDA’s performance. 5) We also present PPL as an auxiliary metric to measure the generation perplexity. Our method outperforms the others due to the high quality of data generation. We also provide experimental comparisons with other DA approaches (SUB² and VDA) and generation speed results in the supplementary file. In conclusion, EPiDA is a powerful and efficient technique.

Method	TREC		AGNews		Sentiment		Irony		Offense	
	1%	10%	0.05%	0.1%	1%	10%	1%	10%	0.1%	1%
CNN	0.722	0.806	0.745	0.826	0.446	0.584	0.534	0.616	0.479	0.548
+EPiDA with EDA	0.745	0.814	0.806	0.829	0.520	0.598	0.579	0.621	0.493	0.551
+EPiDA with CWE	0.737	0.817	0.806	0.826	0.524	0.598	0.578	0.623	0.488	0.556
+EPiDA with TextAttack	0.723	0.838	0.819	0.829	0.527	0.600	0.568	0.631	0.481	0.549
BERT	0.769	0.914	0.759	0.820	0.507	0.602	0.589	0.635	0.508	0.630
+EPiDA with EDA	0.786	0.931	0.813	0.832	0.538	0.641	0.598	0.652	0.525	0.629
+EPiDA with CWE	0.780	0.922	0.821	0.834	0.546	0.642	0.597	0.655	0.517	0.630
+EPiDA with TextAttack	0.762	0.930	0.816	0.839	0.551	0.621	0.595	0.634	0.505	0.626
XLNet	0.746	0.904	0.749	0.776	0.563	0.620	0.576	0.651	0.522	0.626
+EPiDA with EDA	0.756	0.906	0.768	0.790	0.556	0.609	0.588	0.651	0.536	0.625
+EPiDA with CWE	0.750	0.894	0.779	0.795	0.554	0.608	0.592	0.659	0.532	0.627
+EPiDA with TextAttack	0.758	0.909	0.790	0.798	0.555	0.591	0.570	0.654	0.528	0.618

Table 4: Performance comparison with different DA algorithms and classification models on five classification tasks.

4.3 Performance with Different DA Algorithms and Classifiers

EPiDA is a plug-in component that can work with different DA algorithms and classifiers. Here, to check how EPiDA performs with different DA algorithms and classifiers, we consider three frequently-used DA algorithms: rule-based EDA (Wei and Zou, 2019), model-based CWE (Kobayashi, 2018) and Attack-based TextAttack (Morris et al., 2020), and three different classifiers: CNN (Kim, 2014), BERT (Devlin et al., 2019) and XLNet (Yang et al., 2019). And to show that EPiDA can cope with different NLP classification tasks, we present the results on five different tasks: **TREC**, **AGNews**, **Sentiment**, **Irony** and **Offense**. In order to fully evaluate the performance of DA algorithms, we use only a small part of the training data. The experimental results are presented in Tab. 4. Here, we remove the restriction of only one time augmentation so that EPiDA can continuously generate qualified samples. We call this *online augmentation*, to differentiate it from one-time augmentation. As shown in Tab. 4, EPiDA is applicable to various NLP classification tasks. Although these tasks have different forms of data (questions or tweets) and different degrees of classification difficulty, EPiDA boosts performance on these tasks in almost all cases. More details on how EPiDA controls the generation quality are discussed in ablation study. Besides, we can also see that EPiDA works well with the three DA algorithms EDA, CWE and TextAttack. All achieve improved performance in most

ID	DA	REM	CEM	OA	PT	TREC 1%	Irony 1%
1	-	-	-	-	-	0.722	0.534
2	✓	-	-	-	-	0.736	0.474
3	✓	-	-	✓	-	0.723	0.550
4	✓	✓	-	-	✓	0.729	0.557
5	✓	-	✓	-	✓	0.723	0.559
6	✓	✓	✓	-	-	0.734	0.548
7	✓	✓	✓	✓	-	0.739	0.575
8	✓	✓	✓	-	✓	0.740	0.576
9	✓	✓	✓	✓	✓	0.745	0.579

Table 5: Ablation study on TREC 1% and Irony 1%. Here, DA, REM, CEM, OA and PT indicate whether applying DA, REM, CEM, online DA (performing multiple DA operations) and pre-training, respectively.

cases. For the three different classification models: CNN, BERT and XLNet, with the help of EPiDA, they all but XLNet on Sentiment get classification performance improvement, which shows that EPiDA is insensitive to classification models.

4.4 Ablation Study

Here we conduct ablation study to check the effectiveness of different EPiDA configurations. We take CNN as the classifier, EDA as the DA algorithm and report the Macro-F1 score over five repeated experiments on **TREC 1%** and **Irony 1%**. Tab. 5 shows the experimental results.

Effect of REM and CEM. The 4th and 5th rows show the results with only REM and CEM, respectively. Both of them perform better than the baseline (1st row), but not as good as the combined case

(the 8th row). On **TREC** (relatively simple task), REM outperforms CEM (0.729 vs. 0.723), while on **Irony** (relatively hard task), CEM outperforms REM (0.559 vs. 0.557). Using only REM can limitedly boost performance since REM promotes the generation of high diversity samples, which may have wrong labels. And using only CEM is also not enough to fully tap the performance as CEM tends to generate redundant samples.

We also compare our ‘REM + CEM’ with ‘PPL + cosine similarity’ used in LDA (Zuo et al., 2021). Our method achieves the performance of 0.740 and 0.576 on **TREC 1%** and **Irony 1%**, while the latter achieves 0.730 and 0.562. This shows that our ‘REM + CEM’ is more effective.

Effect of online augmentation. Comparing the results of the 2nd and the 3rd rows, the 6th and the 7th rows, the 8th and the 9th rows, we can see that generally online augmentation can boost performance, as online augmentation can generate sufficient qualified samples to train the model.

Effect of pre-training. As REM and CEM use the feedback of the classifier, a pre-trained classification model should be beneficial to REM and CEM. By comparing the results of the 6th and the 8th rows, the 7th and the 9th rows, it is obvious that pre-training can improve performance.

Effect of normalization. In Alg. 1, we normalize s_{div} and s_{qua} . Here, we check the effect of normalization. With the same experimental settings, the performance results on **TREC 1%** and **Irony 1%** without normalization are 0.732 and 0.568, lower than the normalized results 0.740 and 0.576. This shows that normalization is effective.

How to combine REM and CEM? How to combine REM and CEM is actually how to combine the values of s_{div} and s_{qua} . We consider three simple schemes: addition ($s_{tot} = s_{div} + s_{qua}$), multiplication ($s_{tot} = s_{div} * s_{qua}$) and weighted addition ($s_{tot} = \alpha s_{div} + (1 - \alpha) s_{qua}$, α is a hyperparameter to tradeoff REM and CEM). Note that for multiplication, there is possibly an extreme situation: after normalization, s_{div} or s_{qua} may be very small and even approaches 0, then the multiplication result is very small or even zero, which means that REM and CEM do not take effect in sample generation. In our experiments, the multiplication scheme achieves performance of 0.725 and 0.572 on **TREC 1%** and **Irony 1%**, lower than the addition scheme 0.740 and 0.576. As for weighted addition, we find that setting $\alpha = 0.5$ can achieve satisfactory re-

Metric	EDA	REM only	CEM only	EPiDA
Error Rate↓	3.05%	6.75%	0.64%	1.53%
Distance↑($\times 10^{-2}$)	0.54	1.21	0.25	0.78

Table 6: Error rate and distance on **Sentiment** dataset.

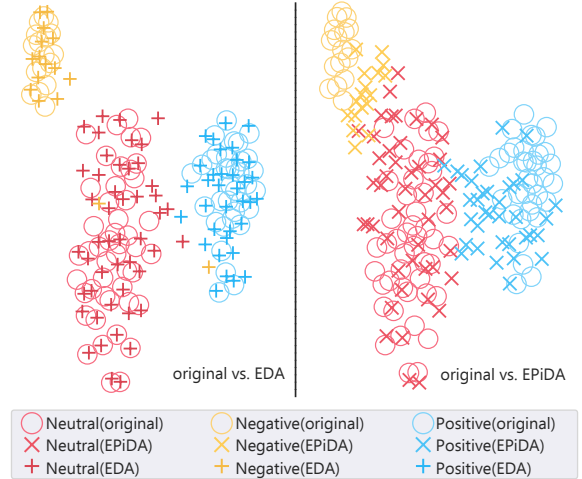


Figure 2: Hidden state visualization of the sentiment analysis task by t-SNE. Left: original and EDA augmented data; Right: original and EPiDA augmented data. Different colors mean different classes (Neutral, Negative and Positive), and different shapes mean different samples (original, generated by EPiDA and EDA).

sults (see the supplementary file). This is actually equal to the addition scheme. Therefore, in our experiments, we use only the addition scheme.

Quality and diversity metrics. Here, we provide another two metrics to verify EPiDA from the perspective of quality and diversity. For quality, we use the augmentation error rate. As for diversity, we calculate the average distance of samples before and after augmentation (ignoring wrong samples). From the perspective of quality and diversity, a good DA should have a *small* error rate but a *large* distance. Experimental results are given in Tab. 6. We can see that EPiDA gets better trade-off between error rate and distance.

Effect of the amplification factor K . The amplification factor K determines the size Km of candidate samples from which m samples are chosen. On the one hand, with a large K , we have more choices, which seems beneficial to diversity. On the other hand, more candidate samples make the selected samples more homogenous, not good for diversity. By grid search, we set K to 3 in our experiments, the experimental results are shown in the supplementary file.

4.5 Visualization Effect of EPiDA

Above we give comprehensive quantitative performance evaluation of EPiDA, here to intuitively illustrate the effectiveness of EPiDA, we visualize some augmented samples of EPiDA, and compare them with that of EDA. Specifically, we utilize BERT as the classifier and visualize its hidden state on the sentiment analysis task via t-SNE (Van der Maaten and Hinton, 2008). Fig. 2 shows the results. In terms of data quality, we find that two negative samples generated by EDA are located in Neural and Positive classes, while samples generated by EPiDA are generally properly located. And in the point of view of diversity, samples generated by EPiDA extend the distributed areas of the original data, while samples generated by EDA are mainly located in the areas of the original samples. This shows that samples generated by EPiDA are more diverse than those generated by EDA.

5 Conclusion

In this paper, we present an easy plug-in data augmentation technique EPiDA to control augmented data diversity and quality via two mechanisms: *relative entropy maximization* and *conditional entropy minimization*. Through extensive experiments, we show that EPiDA outperforms existing methods, and can work well with different DA algorithms and classification models. EPiDA is general, effective, efficient, and easy-to-deploy. In the future, more verification of our method is expected to be conducted on other classification tasks.

Acknowledgement

This work was partially supported by National Key R&D Program of China under grant No. 2021YFC3340302, and Alibaba Innovative Research (AIR) programme under contract No. SCCW802020046613.

References

- Ateret Anaby-Tavor, Boaz Carmeli, Esther Goldbraich, Amir Kantor, George Kour, Segev Shlomov, Naama Tepper, and Naama Zwerdling. 2020. Do not have enough data? deep learning to the rescue! In *Proceedings of the AAI Conference on Artificial Intelligence*, volume 34, pages 7383–7390.
- Jordan T Ash, Chicheng Zhang, Akshay Krishnamurthy, John Langford, and Alekh Agarwal. 2019. Deep batch active learning by diverse, uncertain gradient lower bounds. *arXiv preprint arXiv:1906.03671*.
- Ken Chatfield, Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. 2014. Return of the devil in the details: Delving deep into convolutional nets. *arXiv preprint arXiv:1405.3531*.
- Jiaao Chen, Zichao Yang, and Diyi Yang. 2020. Mixtext: Linguistically-informed interpolation of hidden space for semi-supervised text classification. *arXiv preprint arXiv:2004.12239*.
- Claude Coulombe. 2018. Text data augmentation made simple by leveraging nlp cloud apis. *arXiv preprint arXiv:1812.04718*.
- Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. 2019. Autoaugment: Learning augmentation strategies from data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 113–123.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186.
- Xin Luna Dong, Yaxin Zhu, Zuohui Fu, Dongkuan Xu, and Gerard de Melo. 2021. Data augmentation with adversarial training for cross-lingual nli. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5158–5167.
- Steven Y Feng, Varun Gangal, Jason Wei, Sarath Chandar, Soroush Vosoughi, Teruko Mitamura, and Edward Hovy. 2021. A survey of data augmentation approaches for nlp. *arXiv preprint arXiv:2105.03075*.
- Antigoni Maria Founta, Constantinos Djouvas, Despoina Chatzakou, Ilias Leontiadis, Jeremy Blackburn, Gianluca Stringhini, Athena Vakali, Michael Sirivianos, and Nicolas Kourtellis. 2018. Large scale crowdsourcing and characterization of twitter abusive behavior. In *Twelfth International AAAI Conference on Web and Social Media*.
- Fei Gao, Jinhua Zhu, Lijun Wu, Yingce Xia, Tao Qin, Xueqi Cheng, Wengang Zhou, and Tie-Yan Liu. 2019. Soft contextual data augmentation for neural machine translation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5539–5544.
- Akbar Karimi, Leonardo Rossi, and Andrea Prati. 2021. Aeda: An easier data augmentation technique for text classification. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2748–2754.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1746–1751.

- Sosuke Kobayashi. 2018. Contextual augmentation: Data augmentation by words with paradigmatic relations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 452–457.
- Oleksandr Kolomiyets, Steven Bethard, and Marie Francine Moens. 2011. Model-portability experiments for textual temporal analysis. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 271–276.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25:1097–1105.
- Ashutosh Kumar, Satwik Bhattamishra, Manik Bhandari, and Partha Talukdar. 2019. Submodular optimization-based diverse paraphrasing and its effectiveness in data augmentation. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3609–3619.
- Xin Li and Dan Roth. 2002. Learning question classifiers. In *COLING 2002: The 19th International Conference on Computational Linguistics*.
- Sungbin Lim, Ildoo Kim, Taesup Kim, Chiheon Kim, and Sungwoong Kim. 2019. Fast autoaugment. *arXiv preprint arXiv:1905.00397*.
- Ruibo Liu, Guangxuan Xu, Chenyan Jia, Weicheng Ma, Lili Wang, and Soroush Vosoughi. 2020. Data boost: Text data augmentation through reinforcement learning guided conditional generation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9031–9041.
- Canjie Luo, Yuanzhi Zhu, Lianwen Jin, and Yongpan Wang. 2020. Learn to augment: Joint data augmentation and network optimization for text recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13746–13755.
- John X Morris, Eli Lifland, Jin Yong Yoo, Jake Grigsby, Di Jin, and Yanjun Qi. 2020. Textattack: A framework for adversarial attacks, data augmentation, and adversarial training in nlp. *Proceedings of the 2020 EMNLP*.
- Nathan Ng, Kyunghyun Cho, and Marzyeh Ghassemi. 2020. Ssmba: Self-supervised manifold based data augmentation for improving out-of-domain robustness. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1268–1283.
- Yuyang Nie, Yuanhe Tian, Xiang Wan, Yan Song, and Bo Dai. 2020. Named entity recognition for social media texts with semantic augmentation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1383–1391.
- Tong Niu and Mohit Bansal. 2018. Adversarial oversensitivity and over-stability strategies for dialogue models. In *Proceedings of the 22nd Conference on Computational Natural Language Learning*, pages 486–496.
- Husam Quteineh, Spyridon Samothrakis, and Richard Sutcliffe. 2020. Textual data augmentation for efficient active learning on tiny datasets. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7400–7410.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Sara Rosenthal, Noura Farra, and Preslav Nakov. 2017. Semeval-2017 task 4: Sentiment analysis in twitter. In *Proceedings of the 11th international workshop on semantic evaluation (SemEval-2017)*, pages 502–518.
- Gözde Gül Şahin and Mark Steedman. 2018. Data augmentation via dependency tree morphing for low-resource languages. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5004–5009.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Improving neural machine translation models with monolingual data. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 86–96.
- Haoyue Shi, Karen Livescu, and Kevin Gimpel. 2021. Substructure substitution: Structured data augmentation for nlp. *arXiv preprint arXiv:2101.00411*.
- Connor Shorten and Taghi M Khoshgoftaar. 2019. A survey on image data augmentation for deep learning. *Journal of Big Data*, 6(1):1–48.
- Chenglei Si, Zhengyan Zhang, Fanchao Qi, Zhiyuan Liu, Yasheng Wang, Qun Liu, and Maosong Sun. 2021. Better robustness by more coverage: Adversarial and mixup data augmentation for robust finetuning. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 1569–1576.
- Walter Simoncini and Gerasimos Spanakis. 2021. Seqat-tack: On adversarial attacks for named entity recognition. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 308–318.

- Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. 2015. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9.
- Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-sne. *Journal of machine learning research*, 9(11).
- Cynthia Van Hee, Els Lefever, and Véronique Hoste. 2018. Semeval-2018 task 3: Irony detection in english tweets. In *Proceedings of The 12th International Workshop on Semantic Evaluation*, pages 39–50.
- Vikas Verma, Alex Lamb, Christopher Beckham, Amir Najafi, Ioannis Mitliagkas, David Lopez-Paz, and Yoshua Bengio. 2019. Manifold mixup: Better representations by interpolating hidden states. In *International Conference on Machine Learning*, pages 6438–6447. PMLR.
- William Yang Wang and Diyi Yang. 2015. That’s so annoying!!!: A lexical and frame-semantic embedding based data augmentation approach to automatic categorization of annoying behaviors using# petpeeve tweets. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2557–2563.
- Jason Wei, Chengyu Huang, Shiqi Xu, and Soroush Vosoughi. 2021. Text augmentation in a multi-task view. *arXiv preprint arXiv:2101.05469*.
- Jason Wei and Kai Zou. 2019. Eda: Easy data augmentation techniques for boosting performance on text classification tasks. *arXiv preprint arXiv:1901.11196*.
- Ziang Xie, Sida I Wang, Jiwei Li, Daniel Lévy, Aiming Nie, Dan Jurafsky, and Andrew Y Ng. 2017. Data noising as smoothing in neural network language models. *arXiv preprint arXiv:1703.02573*.
- Yiben Yang, Chaitanya Malaviya, Jared Fernandez, Swabha Swayamdipta, Ronan Le Bras, Ji-Ping Wang, Chandra Bhagavatula, Yejin Choi, and Doug Downey. 2020. G-daug: Generative data augmentation for commonsense reasoning. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*, pages 1008–1025.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Ruslan Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. *arXiv preprint arXiv:1906.08237*.
- Adams Wei Yu, David Dohan, Minh-Thang Luong, Rui Zhao, Kai Chen, Mohammad Norouzi, and Quoc V Le. 2018. Qanet: Combining local convolution with global self-attention for reading comprehension. *arXiv preprint arXiv:1804.09541*.
- Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. 2019. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6023–6032.
- Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. 2017. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*.
- Lu Zhang, Jiandong Ding, Yi Xu, Yingyao Liu, and Shuigeng Zhou. 2021. Weakly-supervised text classification based on keyword graph. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 2803–2813.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. *Advances in neural information processing systems*, 28:649–657.
- Kun Zhou, Wayne Xin Zhao, Sirui Wang, Fuzheng Zhang, Wei Wu, and Ji-Rong Wen. 2021. Virtual data augmentation: A robust and general framework for fine-tuning pre-trained models. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3875–3887.
- Xinyu Zuo, Pengfei Cao, Yubo Chen, Kang Liu, Jun Zhao, Weihua Peng, and Yuguang Chen. 2021. Learnda: Learnable knowledge-guided data augmentation for event causality identification. *arXiv preprint arXiv:2106.01649*.