

Mere Contrastive Learning for Cross-Domain Sentiment Analysis

Yun Luo^{1,2}, Fang Guo², Zihan Liu², Yue Zhang^{2,3}

¹ School of Computer Science And Technology, Zhejiang University, Hangzhou 310024, P.R. China.

² School of Engineering, Westlake University, Hangzhou, China.

³ Institute of Advanced Technology, Westlake Institute for Advanced Study, Hangzhou, China.

{luoyun, guofang, liuzihan, zhangyue}@westlake.edu.cn

Abstract

Cross-domain sentiment analysis aims to predict the sentiment of texts in the target domain using the model trained on the source domain to cope with the scarcity of labeled data. Previous studies are mostly cross-entropy-based methods for the task, which suffer from instability and poor generalization. In this paper, we explore contrastive learning on the cross-domain sentiment analysis task. We propose a modified contrastive objective with in-batch negative samples so that the sentence representations from the same class can be pushed close while those from the different classes become further apart in the latent space. Experiments on two widely used datasets show that our model can achieve state-of-the-art performance in both cross-domain and multi-domain sentiment analysis tasks. Meanwhile, visualizations demonstrate the effectiveness of transferring knowledge learned in the source domain to the target domain and the adversarial test verifies the robustness of our model.

1 Introduction

Sentiment classification (Liu, 2012) has been widely studied by both industry and academia (Blitzer et al., 2007; Li et al., 2013; Yu and Jiang, 2016). For example, the sentiment is positive towards the text ‘*The book is exactly as pictured/described. Cute design and good quality.*’. Early methods rely on labeled data to train models on a specific domain (e.g. DVD reviews, book reviews, and so on), which are labor-intensive and time-consuming (Socher et al., 2013). To address this issue, cross-domain sentiment analysis attracts increasing attention.

Various neural models have been proposed for cross-domain sentiment analysis in recent years (Blitzer et al., 2007; Li et al., 2013; Yu and Jiang, 2016; Zhang et al., 2019; Zhou et al., 2020a). Most methods focus on making the model unable to distinguish the data from which domain by adversarial

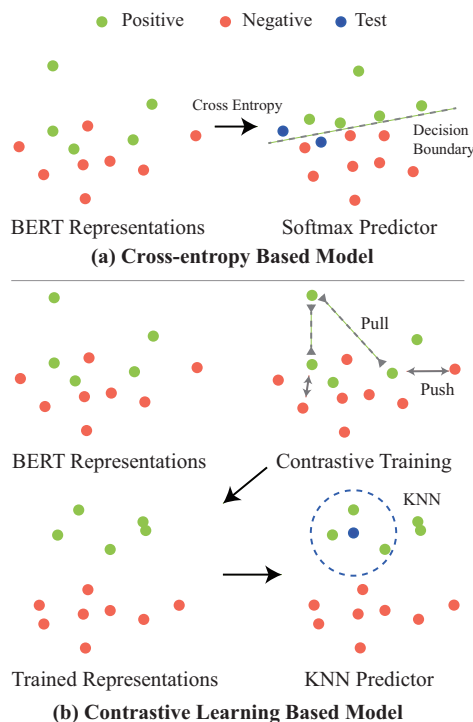


Figure 1: The architectures for cross-entropy-based model and the contrastive-learning-based model.

training or data argumentation, in order to transfer knowledge from source domains to target domains (Du et al., 2020; Liu et al., 2017; Qu et al., 2019; Yang et al., 2022) and some attempt to learn domain-specific knowledge (Zhou et al., 2020a; Liu et al., 2018; Wang et al., 2019). Pre-trained language models (Kenton et al., 2019; Radford et al., 2019; Lewis et al., 2020) have achieved stronger performance compared with previous random initialized models such as LSTM (Long Short-term Memory) in cross-domain tasks. The state-of-the-art models on cross-domain sentiment analysis, such as BERT-DAAT (Du et al., 2020), use unlabeled data to continually train the pre-trained model BERT to transfer knowledge besides adversarial training.

In the representation aspect for the cross-domain

sentiment analysis, there are two key requirements for the representations of sentences: (1) sentence representations in the same domain with the different/the same sentiments should be far from/close to each other; (2) sentence representations of different domains with the same labels should be close. Existing methods are mostly softmax-based method by optimizing cross-entropy loss to achieve the requirements (illustrate in Figure 1 (a)), which suffers from instability across different runs (Zhang et al., 2020; Dodge et al., 2020), poor generalization performance (Liu et al., 2016; Cao et al., 2019), reduction of prediction diversity (Cui et al., 2020), lack of robustness to noisy labels (Zhang and Sabuncu, 2018; Sukhbaatar et al., 2015), or adversarial examples (Elsayed et al., 2018; Nar et al., 2019), especially when supervised data are limited in the cross-domain settings.

To address the above shortcomings, we explore the effectiveness of contrastive learning on the task. Contrastive learning is a similarity-based training strategy, which aims to push the representations from the same class close and those from the different class further apart (Chen et al., 2020; Gao et al., 2021; Neelakantan et al., 2022; Gao et al., 2021). Contrastive learning has been shown effective in solving the problem of anisotropy (Gao et al., 2019), and it has good generalization and robustness (Li et al., 2021; Gao et al., 2021; Gunel et al., 2020; Khosla et al., 2020). Previous work relies mostly on pre-training for representations (Chen et al., 2020; Neelakantan et al., 2022) or multi-task training for semantic textual similarity (Gao et al., 2021), classification (Li et al., 2021; Gunel et al., 2020) and so on, but little work uses mere contrastive learning for supervised tasks. Intuitively, the optimization of contrastive learning is effective in satisfying the requirements of cross-domain sentiment analysis.

We explore **CO**ntrastive learning on **BERT** (COBE) by a modified contrastive loss function with the in-batch negative method on cross-domain sentiment analysis tasks. In the mini-batch, the samples with the same labels are treated as positive pairs, and those with different labels are treated as negative pairs. As shown in Figure 1, the optimization procedure aims to tighten the cluster of samples with the same labels, and push away samples with different labels. After training, the representations of training data and their labels are saved offline as a knowledge base for classifica-

tion. When evaluating the model, a kNN (k-Nearest Neighbors) predictor is used to predict the sentiment of test data, i.e. we search for the k data with the largest cosine similarity in the knowledge base and vote for the final prediction using their labels.

Experiments on two widely used datasets (the cross-domain Amazon dataset (Blitzer et al., 2007) and FDU-MTL (Liu et al., 2017)) show that our model can achieve the state-of-the-art performance in both the cross-domain setting and the multi-domain setting sentiment classification. Visualizations also demonstrate the effectiveness of transferring knowledge learned in the source domain to the target domain. To our knowledge, we are the first to show that contrastive learning outperforms cross-entropy-based models on cross-domain sentiment analysis for both performance and robustness. The code has been released in <https://github.com/LuoXiaoHeics/COBE>.

2 Related Work

Cross-domain sentiment analysis. Due to the heavy cost of obtaining large quantities of labeled data for each domain, many approaches have been proposed for cross-domain sentiment analysis (Blitzer et al., 2007; Li et al., 2013; Yu and Jiang, 2016; Zhang et al., 2019; Zhou et al., 2020a). Ziser and Reichart (2018) and Li et al. (2018a) propose to capture the pivots that are useful for both source domains and target domains. Ganin et al. (2016) propose to use adversarial training with a domain discriminator to learn domain-invariant information, which is one type of solutions for the cross-domain sentiment analysis task (Du et al., 2020; Liu et al., 2017; Qu et al., 2019; Zhou et al., 2020a). These adversarial training methods try to confuse the model unable to classify the data from which domain, transferring the knowledge from source domains to target domains. Besides, Liu et al. (2018) and Cai and Wan (2019) attempt to learn domain-specific information for the different sentiment expressions on different domains. However, these studies rely on minimizing the cross-entropy loss, resulting in the issue of unstable fine-tuning and poor generalization (Gunel et al., 2020; Li et al., 2021; Zhang et al., 2020; Dodge et al., 2020).

Contrastive Learning. Contrastive learning has been widely used in unsupervised learning (Chen et al., 2020; Jing et al., 2021; Wang and Isola, 2020; Khosla et al., 2020; Gao et al., 2021; Neelakantan et al., 2022). Radford et al. (2019) propose

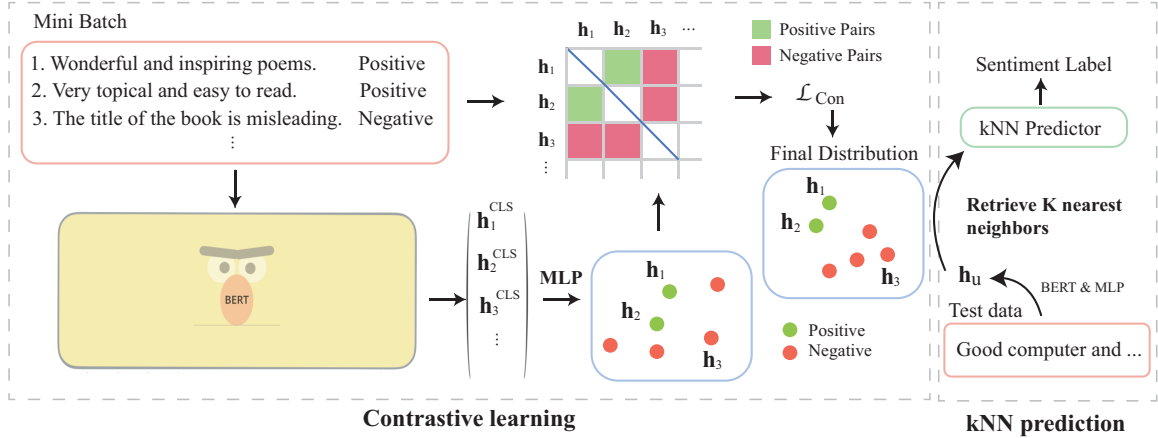


Figure 2: The framework of our contrastive learning for cross-domain sentiment analysis.

to use contrastive learning to learn the representations of both text and images through raw data in unsupervised method, which achieves strong performance on zero-shot task. Neelakantan et al. (2022) propose to use contrastive learning to obtain sentence and code representations and achieve strong performance on downstream tasks such as sentence classification and text search. Wang and Isola (2020) further identify the key properties for contrastive learning as (1) alignment (closeness) of features from positive pairs and (2) uniformity of induced distribution of representations. Gao et al. (2021) uses contrastive learning to learn the sentence representations and theoretically prove that contrastive learning can solve the anisotropy problem (the learned embeddings occupy a narrow cone in the vector space), which limits the expressiveness of representations. It also achieves better results for the semantic textual similarity task using a supervised dataset of natural language inference. Our model differs from the above studies in that we consider contrastive learning in supervised tasks, which uses golden labels to obtain positive/negative pairs for training.

Recently, some studies attempt to incorporate contrastive learning into cross-entropy-based methods by adding InfoNCE loss (Gunel et al., 2020; Li et al., 2021), which aims to solve the shortcomings of cross-entropy loss. Gunel et al. (2020) propose a new SCL loss based on InfoNCE loss to boost the stability and robustness of fine-tuning pre-trained language models. Subsequently, Li et al. (2021) attempt to incorporate kNN predictors to enhance the generalization of prediction in few-shot tasks, using both cross-entropy loss and SCL loss. The above work is similar to ours in making use of contrastive

loss for classification. However, the difference is that we do not use a standard cross-entropy loss, but rely solely on vector space similarity losses for achieving cross-domain classification. To our knowledge, we are the first to conduct sentiment classification without using a cross-entropy loss in natural language processing.

3 Method

Formally, the training data consists of $\{(S_i, Y_i)\}_{i=1}^N$, where $S_i = [s_1, s_2, \dots, s_l]$ is a set of review text, and $Y_i \in \{0, 1\}$ are the corresponding sentiment labels. The model framework is shown in Figure 2. To present our model – **CO**ntrastive learning on **B**ERT (COBE), We first introduce the prediction of sentiment labels using representations based on kNN (Section 3.1), and then describe the training objective to obtain effective representations using contrastive learning (Section 3.2). For comparison, we also describe the standard cross-entropy baseline, named BERT-CE, and a version that adopts adversarial training, named BERT-adv.

3.1 Model

We concatenate the review text S_i with special tokens $[CLS]$ and $[SEP]$ as our model input $X_i = [CLS] S_i [SEP]$, which is fed into BERT model to obtain the hidden states. The hidden state of $[CLS]$ from the last layer of BERT is considered as the representation of the input sequence:

$$\mathbf{h}_i^{CLS} = BERT(X_i)[CLS] \quad (1)$$

BERT-CE and BERT-adv baselines: After obtaining the sentence representation \mathbf{h}_i^{CLS} of input X_i , an MLP (Multi Layer Perceptron) layer project

it to the label space and a softmax layer is adopted to calculate the probability distribution on the labels:

$$p_{ce} = \text{Softmax}(MLP(\mathbf{h}_i^{CLS})). \quad (2)$$

The label with the largest probability is adopted as the prediction result.

COBE: Our model uses the same representation of Eq(1), but adopts a kNN predictor to classify the labels. An MLP layer is then adopted for the dimension reduction:

$$\mathbf{h}_i = MLP(\mathbf{h}_i^{CLS}). \quad (3)$$

To predict the sentiment label of a review text S_u , we calculate the cosine similarity of the sentence representation \mathbf{h}_u with the sentence representations of the training data:

$$\text{sim}(\mathbf{h}_u, \mathbf{h}_i) = \frac{\mathbf{h}_u \cdot \mathbf{h}_i}{\|\mathbf{h}_u\| \cdot \|\mathbf{h}_i\|}, \quad (4)$$

where \mathbf{h}_i is the sentence representation of training data X_i .

We retrieve the k training data whose cosine similarity with \mathbf{h}_u are the largest. We denote the k nearest neighbors as $(\mathbf{h}_i, Y_i) \in \mathcal{K}_u$. The retrieved set is converted to a probability distribution over the labels by applying a softmax with temperature T to the similarity. Using the temperature $T > 1$ can flatten the distribution, and prevent over-fitting to the most similar searches (Khandelwal et al., 2020). The probability distribution on the labels can be expressed as follows:

$$p_k(Y'_u) \propto \sum_{(\mathbf{h}_i, Y_i) \in \mathcal{K}_u} \mathbb{1}_{Y'_u=Y_i} \cdot \exp\left(\frac{\text{sim}(\mathbf{h}_u, \mathbf{h}_i)}{T}\right). \quad (5)$$

The label with the largest probability is regarded as the prediction result.

3.2 Training Objective

BERT-CE baseline: For the cross-entropy-based model, multi-label cross-entropy loss is adopted to optimize the model, which is formulated as follows:

$$\mathcal{L}_{cls} = -\frac{1}{M} \sum_{(X_i, Y_i)} Y_i \log p_{ce}(Y_i), \quad (6)$$

BERT-adv baseline: Besides the cross-entropy loss, BERT-adv adds a domain discriminator (Du et al., 2020; Ganin et al., 2016) to the standard

model and adopts adversarial training to transfer knowledge from source domains to target domains.

Given the sentence and its domain label (X_i, D_i) , the representation \mathbf{h}_i^{CLS} obtained in Eq(1) goes through an additional gradient reversal layer (GRL) (Ganin et al., 2016), which can be denoted as a ‘pseudo-function’ $D_\lambda(x)$. The GRL reverses the gradient by applying a negative scalar λ . The forward- and backward- behaviors can be described:

$$D_\lambda(x) = x, \quad \frac{\partial D_\lambda(x)}{\partial x} = -\lambda I, \quad (7)$$

where λ is a hyper-parameter and I is the gradients calculated on \mathbf{h}_i^{CLS} (but it is multiplied with $-\lambda$ to back-propagate). Then a linear layer project \mathbf{h}_i^{CLS} to the label space and a softmax layer is adopted to calculate the distribution on domain labels:

$$p_d = \text{Softmax}(W_d \mathbf{h}_i^{CLS} + b_d), \quad (8)$$

where W_d and b_d are the learnable parameters. The training target is to minimize the cross-entropy for all data from the source and target domains (note that the data from target domains are unlabeled on sentiment) in order to make the model unable to predict the domain labels:

$$\mathcal{L}_{dom} = -\frac{1}{M} \sum_{(X_i, D_i)} D_i \log p_d(D_i). \quad (9)$$

For BERT-adv, the training loss of sentiment classification (Eq.7) and domain classification (Eq.10) are jointly optimized:

$$\mathcal{L}_{adv} = \mathcal{L}_{cls} + \mathcal{L}_{dom} \quad (10)$$

COBE: The baselines adopt \mathcal{L}_{cls} to tighten the representations of the same/different labels close (apart), and adopt \mathcal{L}_{dom} to mix up the representations of different domains with the same label. However, COBE uses a single training objective of contrastive learning to achieve the both goals. We apply in-batch negatives (Yih et al., 2011; Sohn, 2016) to learn sentence representations through contrastive learning, which has been widely used in unsupervised representation learning (Radford et al., 2021; Jia et al., 2021). For each example in the mini-batch of M samples, we treat the other samples with different golden labels as negative pairs, and the samples with the same golden labels as positive pairs. For example in Figure 2, the sentence pair (1,2) is positive pairs, and the sentence

S → T	B→D	B→E	B→K	D→B	D→E	D→K	E→B	E→D	E→K	K→B	K→D	K→E	Avg
DANN	82.30	77.60	76.10	81.70	79.70	77.35	78.55	79.70	83.95	79.25	80.45	86.65	80.29
PBLM	84.20	77.60	82.50	82.50	79.60	83.20	71.40	75.00	87.80	74.20	79.80	87.10	80.40
HATN	86.10	85.70	85.20	86.30	85.60	86.20	81.00	84.00	87.90	83.30	84.50	87.00	85.10
ACAN	83.45	81.20	83.05	82.35	82.80	78.60	79.75	81.75	83.35	80.80	82.10	86.60	82.15
IATN	86.80	86.50	85.90	87.00	86.90	85.80	81.80	84.10	88.70	84.70	84.10	87.60	85.90
BERT-CE	88.96	86.15	89.05	89.40	86.55	87.53	86.50	87.95	91.60	87.55	87.95	90.45	88.25
BERT-CE*	55.40	56.55	54.05	55.10	57.25	53.75	55.50	56.00	55.55	52.30	52.75	54.15	54.86
BERT-adv	89.70	87.30	89.55	89.55	86.05	87.69	87.15	86.05	91.91	87.65	87.72	86.05	88.56
DAAT	89.70	89.57	90.75	90.86	89.30	87.53	88.91	90.13	93.18	87.98	88.81	91.72	90.12
COBE*	82.17	83.65	83.12	79.82	78.87	82.58	75.95	79.53	86.10	78.55	76.95	85.17	80.95
COBE (proposed)	90.05	90.45	92.90	90.98	90.67	92.00	87.90	87.87	93.33	88.38	87.43	92.58	90.39

Table 1: Results on the cross-domain Amazon dataset. BERT-CE* and COBE* refer to the models fixing the parameters of BERT, and only tuning the parameters of MLP layer. (B for the Books domain, D for the DVD domain, E for the Electronics domain, and K for the Kitchen domain, respectively.)

pairs (1,3) and (2,3) are negative pairs. For each review X_i we denote N_i^+ as the set of reviews with the same label of X_i in the mini-batch. Then the contrastive loss function can be defined as follows:

$$\mathcal{L}_{Com} = \sum_i^M -\frac{1}{M} \log \frac{\sum_{k \in N_i^+} \exp(\text{sim}(\mathbf{h}_i, \mathbf{h}_k)/\tau)}{\sum_{i \neq j}^M \exp(\text{sim}(\mathbf{h}_i, \mathbf{h}_j)/\tau)} \quad (11)$$

where τ is a temperature hyper-parameter. The loss function can alleviate the negative effect of the situations where there is no positive pairs for any training instance in the batch.

The usage of in-batch negatives enables re-use of computation both in the forward and backward pass making training highly efficient.

4 Experiments

We conduct experiments on both the cross-domain settings (train models on source domains and test on another one) and the multi-domain settings (train and test models on the same domains). To verify the effectiveness of our model, we also visualize the representations (Section 4.3) and carry out further analysis such as model robustness (Section 4.4).

4.1 Settings

Datasets. We test our contrastive learning method on two widely used datasets, the cross-domain Amazon dataset, and the FDU-MTL dataset. The cross-domain Amazon dataset (Blitzer et al., 2007) contains 4 domains: Books (B), DVD (D), Electronics (E) and Kitchen (K). Each domain contains 2000 Amazon review samples. Following the setting of previous work (Ganin et al., 2016; Ziser and Reichart, 2018; Du et al., 2020), we test the model

on 12 tasks. The model is trained on the source domain data and tested on the target domain data.

Furthermore, we also evaluate our model on FDU-MTL, which is an Amazon reviews dataset with data on 16 domains (Liu et al., 2017). The training set, development set, and test set are split in the original dataset, (the statistics are shown in Appendix A). We carry out experiments on the multi-domain setting, (i.e. train the model on the whole 16 domains, and evaluate the model on the test on the whole 16 domains), and on the 15-1 cross-domain setting (i.e. train the model on 15 domains, and test the model on the 1 domain left).

Baselines. For the cross-domain Amazon dataset, we compare our model with several strong baselines in cross-domain sentiment analysis: DANN (Ganin et al., 2016), PBLM (Ziser and Reichart, 2018), HATN (Li et al., 2018b), IATN (Qu et al., 2019), DAAT (Du et al., 2020), BERT-CE and BERT-CE* (* for fixing the BERT parameters). We adopt the results of baselines reported in Zhou et al. (2020b) and Du et al. (2020). We also adopt BERT-adv as our baselines introduced in Section 3.2.

On FDU-MTL, we compare our model with ASP (Liu et al., 2017), DSR-at (Zheng et al., 2018), DAEA and DAEA-B (DAEA-BERT) (Cai and Wan, 2019). The DAEA-B is regarded as the state-of-the-art model on FDU-MTL (excluding the model SentiX (Zhou et al., 2020b), which uses a large corpus (about 241 million reviews) to continually train BERT for sentiment tasks). Note that, it is unfair that previous studies do not adopt the BERT-CE model for multi-domain experiments for comparison. In this study, we also consider BERT-CE, and BERT-CE* on the multi-domain setting as baselines. For the multi-domain task, the objective of

adversarial training is redundant, thus we mainly compare COBE with baselines BERT-CE.

Implementation Details. We perform experiments using the official pre-trained BERT model provided by Huggingface¹. We train our model on 1 GPU (Nvidia GTX2080Ti) using the Adam optimizer (Kingma and Ba, 2014). For the cross-domain Amazon dataset (FDU-MTL), the max sequence length for BERT is 256 (128), and the batch size M is 8 (32). The max sequence lengths are set in such values for comparison with previous models. The initial learning rate is $2e-5$ ($1e-4$) for BERT-unfixed (BERT-fixed) models, and each model is trained for 20 epochs. The hyper-parameters of temperatures τ is 0.05, T is 5, and the number of nearest neighbors k is 3 (without losing generality, we do not search for the best hyper-parameters through grid-search). Through the training of our model, no development set is applied to find the best checkpoints, but stop until the training step is reached. During the test procedure, we adopt FAISS IndexFlat Index (Johnson and Guestrin, 2018) to accelerate the speed to find the k nearest neighbors. We average the results with 3 different random seeds.

4.2 Results

Results on the Cross-Domain Amazon Dataset.

The results are shown in Table 1. Overall, our model COBE achieves state-of-the-art performances with an accuracy of 90.39% on average for the 12 cross-domain tasks. It achieves state-of-the-art performance in 9 of 12 tasks. The result is 2.14% higher than that of BERT-CE (88.25%), which indicates that our proposed contrastive learning method can be more effective and generalized than methods based on cross-entropy loss. COBE is also 1.83% higher than BERT-adv (88.56%), which implies that directly pushing the representations of different domains with the same (different) labels close (apart) results in a strong performance on the cross-domain sentiment classification.

DAAT uses the unlabeled data from the source domain and the target domain to continually train BERT to mix the information of the source domain and target domain. Then the training objective of cross-entropy and the domain discriminator are jointly optimized to obtain the sentiment classification model. The average accuracy of our model is 0.27% higher than that of DAAT, which uses ad-

ditional data to continually train BERT to transfer knowledge in the source domain to the target domain. Although DAAT achieves great performance, it is more time-consuming and resource-wasting compared with solely using contrastive learning. In the tasks of $E \rightarrow B$, $E \rightarrow D$, and $K \rightarrow D$, the accuracies of our model are smaller compared with DAAT, and the possible reason can be that the source domains' data have less shared information with the target domains. But with unlabeled data for continual training, some domain-specific information is extracted in DAAT and further results in a better performance.

Moreover, the average accuracy of the model COBE* (82.05%) outperforms that of BERT-CE* (54.86%) with a large margin, where the parameters of BERT are fixed (corresponding to the scenario that pre-trained models are too large for fine-tuning). The model BERT-CE* fails to predict the sentiments of the target domain using cross-entropy-based methods, but with contrastive learning, it can obtain strong results (similar performance to BERT-CE). But the performance of models fixing BERT parameters is still largely worse than that of unfixed models.

Results on FDU-MTL. First, we test our model in the multi-domain setting, training the model on the data of 16 domains and evaluating it on the whole test data. The results are shown in Table 2. Our model achieves the state-of-the-art performance with an accuracy of 91.49% on average, and in the 12 of 16 domains, it achieves the state-of-the-art performance. The accuracy is 0.67% higher than that of BERT-CE, and 0.99% higher than that of DAEA-B. In particular, using BERT-CE solely in the multi-domain setting can achieve competitive performance (90.82%), which is neglected by previous studies. The accuracy of our model COBE on the IMDB data is lower than DSA with a large margin, which may result from the max sequence length for BERT being 128, much smaller than the average sequence length in IMDB (128 to 256). Our model COBE* achieves an accuracy of 83.35% in the multi-domain setting, which is also higher than that of BERT-CE* with a margin of 3.14%.

Then we also evaluate our model in the 15-1 settings, referring to that we train the model on 15 domains and test it on the domain left (shown in Table 3). Our model achieves state-of-the-art performance with an accuracy of 90.03%. The accuracies in 14 of 16 tasks are larger than that in previous

¹<https://huggingface.co/>

Domain	ASP	DA	DSA	DAEA	DAEA-B	BERT-CE*	BERT-CE	COBE*	COBE
Books	84.00	88.50	89.10	89.00	N/A	81.33	90.67	85.17	90.17
Electronics	86.80	89.00	87.90	91.80	N/A	82.17	91.92	82.92	93.58
DVD	85.50	88.00	88.10	88.30	N/A	78.83	89.00	79.42	89.67
Kitchen	86.20	89.00	85.90	90.30	N/A	79.92	91.17	81.33	91.50
Apparel	87.00	88.80	87.80	89.00	N/A	83.33	92.08	87.25	92.33
Camera	89.20	91.80	90.00	92.00	N/A	81.83	93.25	87.50	93.58
Health	88.20	90.30	92.90	89.80	N/A	81.25	93.33	85.00	93.92
Music	82.50	85.00	84.10	88.00	N/A	79.42	88.92	80.33	90.33
Toys	88.00	89.50	85.90	91.80	N/A	78.25	92.41	83.75	93.42
Video	84.50	89.50	90.30	92.30	N/A	78.17	90.33	83.67	89.91
Baby	88.20	90.50	91.70	92.30	N/A	82.33	93.00	84.42	93.92
Magazines	92.20	92.00	92.10	96.50	N/A	83.41	93.75	89.67	94.08
Software	87.20	90.80	87.00	92.80	N/A	83.42	92.42	85.33	93.42
Sports	85.70	89.80	85.80	90.80	N/A	78.50	91.50	84.50	92.83
IMDB	85.50	89.80	93.80	90.80	N/A	76.43	86.33	76.50	86.91
MR	76.70	75.50	73.30	77.00	N/A	74.75	83.00	76.83	84.33
Avg	86.09	88.61	87.86	90.16	90.50	80.21	90.82	83.35	91.49

Table 2: Results on FDU-MTL in the multi-domain setting. BERT-CE* and COBE* refer to the models fixing the parameters of BERT, and only tuning the parameters of MLP layer.

	ASP	DSR-at	DAEA	COBE
Books	81.50	85.80	87.30	90.67
Electronics	83.80	89.50	85.80	92.33
DVD	84.50	86.30	88.80	87.50
Kitchen	87.50	88.30	88.00	90.75
Apparel	85.30	85.80	88.00	91.16
Camera	85.30	88.80	90.00	91.67
Health	86.00	90.50	91.00	94.33
Music	81.30	84.80	86.50	89.17
Toys	88.00	90.30	90.30	92.33
Video	86.80	85.30	91.30	88.50
Baby	86.50	84.80	90.30	93.17
Magazines	87.00	84.00	88.50	90.50
Software	87.00	90.80	89.80	90.82
Sports	87.00	87.00	90.50	92.15
IMDB	84.00	83.30	85.80	86.58
MR	72.00	76.30	75.50	78.91
Avg	84.59	86.35	87.96	90.03

Table 3: Results on FDU-MTL in the 15-1 setting.

studies. It is 2.07% higher than the average accuracy of DAEA. The experimental results also show that contrastive learning can perform better than cross-entropy-based models with adversarial training for cross-domain sentiment analysis.

4.3 Visualization

The visualization of the sentence representations \mathbf{h}_i in COBE is shown in Figure 3. For the B->K (Books->Kitchen) task in Figure 3 (a), first, the representations of positive and negative data are separated acutely with a large margin between each other. Second, representations of source and target domains with the same labels are close to each other, which means the knowledge learned from the source domain is transferred to the target domain effectively. For the multi-domain setting in Figure

3 (b) (left), we can observe that the representations with the same labels are separated into different clusters w.r.t the labels, and in the sentence representations with the same label but different domains mix up well, which satisfy the requirements of the cross-domain sentiment analysis.

To further compare the contrastive learning method with cross-entropy-based methods, we illustrate the representations of the source domain and the target domain in COBE, BERT-CE in Figure 3 (c)(d), respectively (visualizations of COBE* and BERT-CE* are shown in Appendix). Obviously, the sentence representations are separated in the target domain in BERT-CE less effective than that in COBE. The visualizations show the effectiveness of contrastive learning in transferring the learned knowledge in the source domain to the target domain. Meanwhile, it demonstrates operating the sentence representations in the feature space has a strong generalization ability in the cross-domain sentiment analysis tasks.

4.4 Robustness Analysis

We evaluate our model on adversarial samples generated by using the well-known substitution-based adversarial attack method—Textfooler (Jin et al., 2020). Given an input X_i and a pre-trained classification model F , a valid adversarial sample X_i^{adv} should conform the following requirements:

$$F(X_i) \neq F(X_i^{adv}), \quad Sim(X_i, X_i^{adv}) \geq \epsilon. \quad (12)$$

where Sim is a similarity function and ϵ is the minimum similarity between the original input and the adversarial sample, which is often a semantic

Text	Gold Label	Output
Test Data. This story is true to life living in south west and west phila. It brought back many memories and changing the names did not bother me. I really enjoyed reading about life the way it was back in the 55 to 70 era.	Positive	Positive
k nearest neighbors.		
(1) Have to be honest and say that I haven't seen many independent films, but I thought this one was very well done. The direction and cinematography were engaging without becoming a distraction.	Positive	
(2) I bought this wireless weather station as a gift. The recipient loves it. For the price, he is really enjoying it.	Positive	
(3) I think j-14 is a really good magazine if u like to hear the latest gossip about all your favourite celebrity 's, or if u like to get nice posters of all the hot celebrity 's.	Positive	

Table 4: Case Study on FDU-MTL.

Original Text: DEF. NOT A GOOD TANK. You look at them in a picture frame, the fish are crammed in there.
Adversarial Text: DEF. Not a alright tank. you look at them in a photography sashes , the fish are teeming in there.

Table 5: Adversarial example based on TextFooler.

	BERT-CE*	BERT-CE	COBE*	COBE
Books	42.50	71.50	69.00	78.00
Multi-	49.50	73.50	72.50	81.00

Table 6: Results on the adversarial samples. Books for the trained model using the data of Books domain in the cross-domain Amazon dataset, and Multi- for the trained model using multi-domain data in FDU-MTL.

and syntactic similarity function. The details for generation refer to Jin et al. (2020). An adversarial sample is shown in Table 5, where the sentence semantic information is not corrupted, but some words are replaced.

We test our model (trained using Books data in the cross-domain Amazon dataset) with 200 adversarial samples from the Kitchen domain, and our model (trained using multi-domain data in FDU-MTL) with 200 adversarial samples randomly selected from the multi-domain test data. The results are shown in Table 6. Our model COBE achieves 78.00% and 81.00% accuracies for the two kind of adversarial data, which are 6.5% and 7.5% higher than BERT-CE. Meanwhile, the model COBE* outperforms BERT-CE* with a large margin (26.5% and 23%). The results demonstrate that contrastive-learning based models have better robustness than cross-entropy-based models.

4.5 Case Study

The case study is shown in Table 4. As can be observed, the k nearest neighbors of the test data (Books) are reviews from different domains (Video, Electronics and Magazines) with positive labels,

and it outputs the correct label for the test data. Note that the key sentiment information is similar for the original text and neighbors in the case such as ‘enjoy’, ‘engaging’, ‘enjoying’ and ‘favorite’. It shows that our model can learn effective information from multi-domain data for the sentiment classification task, and the representations of different domains mix up well, which serve as a strong sentiment knowledge base for the classification.

5 Conclusion

We explored the contrastive learning method in the cross-domain sentiment analysis task. We proposed a suitable contrastive loss for the supervised sentiment analysis task with the in-batch negatives method. Experiments on two standard datasets showed the effectiveness of our model. Visualizations also demonstrated the effectiveness of transferring knowledge learned in the source domain to the target domain. We also showed that our model has stronger robustness than cross-entropy-based models through the adversarial test.

Acknowledgements

Yue Zhang is the corresponding author. We would also like to thank the anonymous reviewers for the detailed and thoughtful reviews. The work is funded by the Zhejiang Province Key Project 2022SDXHDX0003.

References

- John Blitzer, Mark Dredze, and Fernando Pereira. 2007. Biographies, bollywood, boom-boxes and blenders: Domain adaptation for sentiment classification. In *Proceedings of the 45th annual meeting of the association of computational linguistics*, pages 440–447.
- Yitao Cai and Xiaojun Wan. 2019. Multi-domain sentiment classification based on domain-aware embedding and attention. In *IJCAI*, pages 4904–4910.

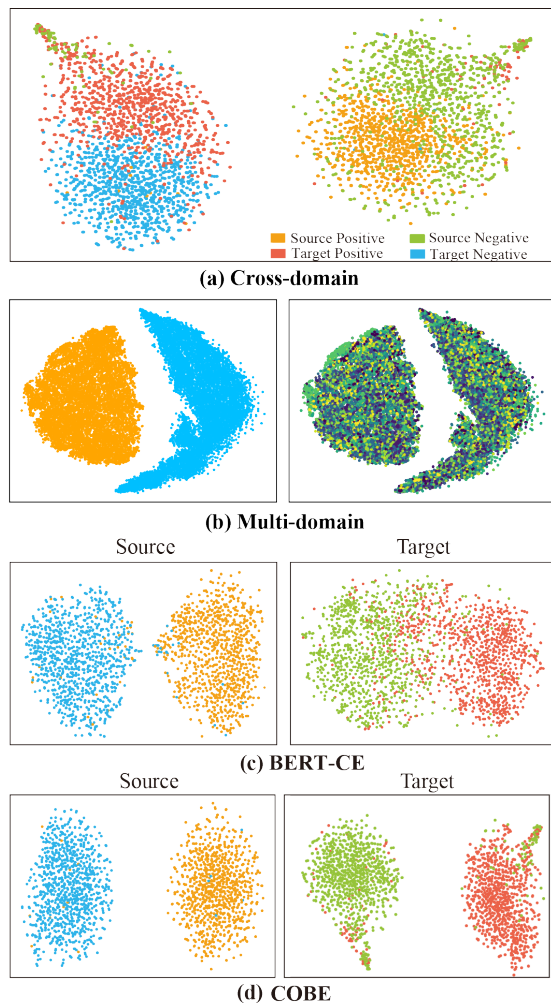


Figure 3: Visualizations of the sentence representations. We use t-SNE to transfer 192-dimensional feature space into two-dimensional space. (a)(c)(d) for the representations of the B->K task in the cross-domain Amazon dataset; (b) for those of the multi-domain task in FDU-MTL (colors in left for sentiment labels and right for domains).

Kaidi Cao, Colin Wei, Adrien Gaidon, Nikos Arachiga, and Tengyu Ma. 2019. Learning imbalanced datasets with label-distribution-aware margin loss. *Advances in neural information processing systems*, 32.

Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A simple framework for contrastive learning of visual representations.

Shuhao Cui, Shuhui Wang, Junbao Zhuo, Liang Li, Qingming Huang, and Qi Tian. 2020. Towards discriminability and diversity: Batch nuclear-norm maximization under label insufficient situations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3941–3950.

Jesse Dodge, Gabriel Ilharco, Roy Schwartz, Ali Farhadi, Hannaneh Hajishirzi, and Noah Smith. 2020.

Fine-tuning pretrained language models: Weight initializations, data orders, and early stopping. arXiv.

Chunning Du, Haifeng Sun, Jingyu Wang, Qi Qi, and Jianxin Liao. 2020. Adversarial and domain-aware bert for cross-domain sentiment analysis. In *Proceedings of the 58th annual meeting of the Association for Computational Linguistics*, pages 4019–4028.

Gamaleldin Elsayed, Dilip Krishnan, Hossein Mobahi, Kevin Regan, and Samy Bengio. 2018. Large margin deep networks for classification. *Advances in neural information processing systems*, 31.

Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. 2016. Domain-adversarial training of neural networks. *The journal of machine learning research*, 17(1):2096–2030.

Jun Gao, Di He, Xu Tan, Tao Qin, Liwei Wang, and Tie Yan Liu. 2019. Representation degeneration problem in training natural language generation models. *7th International Conference on Learning Representations, ICLR 2019*, pages 1–14.

Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. Simcse: Simple contrastive learning of sentence embeddings. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6894–6910.

Beliz Gunel, Jingfei Du, Alexis Conneau, and Veselin Stoyanov. 2020. Supervised contrastive learning for pre-trained language model fine-tuning. In *International Conference on Learning Representations*.

Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. 2021. Scaling up visual and vision-language representation learning with noisy text supervision. In *International Conference on Machine Learning*, pages 4904–4916. PMLR.

Di Jin, Zhijing Jin, Joey Tianyi Zhou, and Peter Szolovits. 2020. Is bert really robust? a strong baseline for natural language attack on text classification and entailment. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 8018–8025.

Li Jing, Pascal Vincent, Yann LeCun, and Yuandong Tian. 2021. Understanding dimensional collapse in contrastive self-supervised learning. *arXiv preprint arXiv:2110.09348*.

Tyler B Johnson and Carlos Guestrin. 2018. Training deep models faster with robust, approximate importance sampling. *Advances in Neural Information Processing Systems*, 31.

Jacob Devlin, Kenton, Chang Ming-Wei, and Lee Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT*, pages 4171–4186.

- Urvashi Khandelwal, Angela Fan, Dan Jurafsky, Luke Zettlemoyer, and Mike Lewis. 2020. Nearest neighbor machine translation. *arXiv preprint arXiv:2010.00710*.
- Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. 2020. Supervised contrastive learning. *Advances in Neural Information Processing Systems*, 33:18661–18673.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Linyang Li, Demin Song, Ruotian Ma, Xipeng Qiu, and Xuanjing Huang. 2021. Knn-bert: Fine-tuning pre-trained models with knn classifier. *arXiv preprint arXiv:2110.02523*.
- Shoushan Li, Yunxia Xue, Zhongqing Wang, and Guodong Zhou. 2013. Active learning for cross-domain sentiment classification. In *Twenty-Third International Joint Conference on Artificial Intelligence*.
- Zheng Li, Ying Wei, Yu Zhang, and Qiang Yang. 2018a. Hierarchical attention transfer network for cross-domain sentiment classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Zheng Li, Ying Wei, Yu Zhang, and Qiang Yang. 2018b. [Hierarchical Attention Transfer Network for Cross-Domain Sentiment Classification](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1).
- Bing Liu. 2012. Sentiment analysis and opinion mining. *Synthesis lectures on human language technologies*, 5(1):1–167.
- Pengfei Liu, Xipeng Qiu, and Xuan-Jing Huang. 2017. Adversarial multi-task learning for text classification. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1–10.
- Qi Liu, Yue Zhang, and Jiangming Liu. 2018. [Learning domain representation for multi-domain sentiment classification](#). *NAACL HLT 2018 - 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies - Proceedings of the Conference*, 1:541–550.
- Weiyang Liu, Yandong Wen, Zhiding Yu, and Meng Yang. 2016. Large-margin softmax loss for convolutional neural networks. In *ICML*, volume 2, page 7.
- Kamil Nar, Orhan Ocal, S Shankar Sastry, and Kannan Ramchandran. 2019. Cross-entropy loss and low-rank features have responsibility for adversarial examples. *arXiv preprint arXiv:1901.08360*.
- Arvind Neelakantan, Tao Xu, Raul Puri, Alec Radford, Jesse Michael Han, Jerry Tworek, Qiming Yuan, Nikolas Tezak, Jong Wook Kim, Chris Hallacy, et al. 2022. Text and code embeddings by contrastive pre-training. *arXiv preprint arXiv:2201.10005*.
- Xiaoye Qu, Zhikang Zou, Yu Cheng, Yang Yang, and Pan Zhou. 2019. Adversarial category alignment network for cross-domain sentiment classification. 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 2496–2508.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 1631–1642.
- Kihyuk Sohn. 2016. Improved deep metric learning with multi-class n-pair loss objective. *Advances in neural information processing systems*, 29.
- Sainbayar Sukhbaatar, Joan Bruna, Manohar Paluri, Lubomir Bourdev, and Rob Fergus. 2015. Training convolutional networks with noisy labels. In *3rd International Conference on Learning Representations, ICLR 2015*.
- Guoyin Wang, Yan Song, Yue Zhang, and Dong Yu. 2019. Learning word embeddings with domain awareness. *arXiv preprint arXiv:1906.03249*.
- Tongzhou Wang and Phillip Isola. 2020. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *International Conference on Machine Learning*, pages 9929–9939. PMLR.

- Linyi Yang, Lifan Yuan, Leyang Cui, Wenyang Gao, and Yue Zhang. 2022. Factmix: Using a few labeled in-domain examples to generalize to cross-domain named entity recognition. *arXiv preprint arXiv:2208.11464*.
- Wen-tau Yih, Kristina Toutanova, John C Platt, and Christopher Meek. 2011. Learning discriminative projections for text similarity measures. In *Proceedings of the fifteenth conference on computational natural language learning*, pages 247–256.
- Jianfei Yu and Jing Jiang. 2016. Learning sentence embeddings with auxiliary tasks for cross-domain sentiment classification. In *Proceedings of the 2016 conference on empirical methods in natural language processing*, pages 236–246.
- Kai Zhang, Hefu Zhang, Qi Liu, Hongke Zhao, Hengshu Zhu, and Enhong Chen. 2019. Interactive attention transfer network for cross-domain sentiment classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 5773–5780.
- Tianyi Zhang, Felix Wu, Arzoo Katiyar, Kilian Q Weinberger, and Yoav Artzi. 2020. Revisiting few-sample bert fine-tuning. *arXiv preprint arXiv:2006.05987*.
- Zhilu Zhang and Mert Sabuncu. 2018. Generalized cross entropy loss for training deep neural networks with noisy labels. *Advances in neural information processing systems*, 31.
- Jun Zhao, Tao Gui, Qi Zhang, and Yaqian Zhou. 2021. A relation-oriented clustering method for open relation extraction. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 9707–9718.
- Renjie Zheng, Junkun Chen, and Xipeng Qiu. 2018. Same representation, different attentions: shareable sentence representation learning from multiple tasks. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, pages 4616–4622.
- Jie Zhou, Junfeng Tian, Rui Wang, Yuanbin Wu, Wenming Xiao, and Liang He. 2020a. Sentix: A sentiment-aware pre-trained model for cross-domain sentiment analysis. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 568–579.
- Jie Zhou, Junfeng Tian, Rui Wang, Yuanbin Wu, Wenming Xiao, and Liang He. 2020b. [SentiX: A sentiment-aware pre-trained model for cross-domain sentiment analysis](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 568–579, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Yftah Ziser and Roi Reichart. 2018. Pivot based language modeling for improved neural domain adaptation. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1241–1251.

A Statistics for FDU-MTL.

Domain	Train	Dev	Test	Avg. Length
Books	1400	200	400	159
Electronics	1398	200	400	101
DVD	1400	200	400	173
Kitchen	1400	200	400	89
Apparel	1400	200	400	57
Camera	1397	200	400	130
Health	1400	200	400	81
Music	1400	200	400	136
Toys	1400	200	400	90
Video	1400	200	400	156
Baby	1300	200	400	104
Magazines	1370	200	400	117
Software	1315	200	400	129
Sports	1400	200	400	94
IMDB	1400	200	400	269
MR	1400	200	400	21

Table 7: Statistics of FDU-MTL.

B Reconstruction Loss

We attempt to reconstruct the representations of BERT which means another MLP layer is applied by $\mathbf{h}_i^{rec} = MLP(\mathbf{h}_i)$. Then a reconstruction loss of MSE (mean-squared loss) is added to retain the semantic information, $\mathcal{L}_{rec} = \|\mathbf{h}_i^{rec} - \mathbf{h}_i^{CLS}\|$ as (Zhao et al., 2021). But little improvement (an average accuracy of 90.81% on the FDU-MTL multi-domain setting and 90.13% on the cross-domain Amazon dataset) is obtained, which is 0.68% and 0.26% lower than COBE, respectively. It indicates that the reconstruction loss is not suitable for the task of cross-domain sentiment analysis.

C SCL Loss

To verify the effectiveness of our propose loss function, we compare our contrastive learning loss with the SCL loss (Gunel et al., 2020), which can be formulated as follows:

$$\mathcal{L}_{SCL} = - \sum_i^M \frac{1}{|N_i^+|} \sum_{k \in N_i^+} \log \frac{\exp(\text{sim}(\mathbf{h}_i, \mathbf{h}_k)/\tau)}{\sum_{i \neq j}^M \exp(\text{sim}(\mathbf{h}_i, \mathbf{h}_j)/\tau)} \quad (13)$$

In fairness, we use the kNN predictor the same as our proposed model. The model with SCL loss achieves an average accuracy of 91.13% on the FDU-MTL multi-domain setting and 90.05% on the cross-domain Amazon dataset (0.36% and 0.34% lower than COBE, respectively). The experiments prove the effectiveness of our proposed loss function with the in-batch negative samples, which aims to tighten all the samples of the same labels as positive pairs. The conclusion is different from




Figure 4: Evaluation with respect to different numbers of k for one random seed.

that in Khosla et al. (2020), whose experiments demonstrate that the separately calculating each positive pair separately (SCL loss) achieves better results for image classification. It may results from the reason that the batch size influences the results of the two methods, and our batch sizes (8 and 32) are comparatively smaller compared with their study (6144), which may motivate further theoretical analysis.

D Influence of k

In order to discover the sensitivity of our model to the influence of k for the kNN predictor (shown in Figure 4), we evaluate our model with respect to the different numbers of k . As observed, the accuracies of COBE stop to increase and keep stable when $k \geq 5$, which indicates that the model is little sensitive with the hyper-parameter k . The phenomenon demonstrates that the sentence representations learned by COBE are effectively separated and stable for classification.

Original Text	Adversarial Text
Very nice iron!. This is a great iron. It's quite heavy, but I like that. It really gets out the wrinkles. I don't even mind ironing any more.	Awfully sweet iron! This is a whopping iron it 's quite heavy, but I like that it really gets out the wrinkles I don't even mind ironing any more.
A good idea, disappointing in use. These silicone pot holders are indeed brightly colored, easy to wash in the dishwasher, and protective even when wet. They are also clumsily stiff at the same time as they are slippery, the net result being a miserable failure in the kitchen. They are useful for protecting a counter from a hot pot, but not for picking the hot pot up.	A good ideas, agonizing in use these silicon pot holders are indeed brightly colour, easy to wash in the dishwasher, and protective even when clammy they are also clumsily painstaking at the same time as they are slippery, the net raison being a miserable failure in the kitchen they are useful for protecting a counter from a hot pot, but not for picking the hot pot up.
Love this piece. I just bought this piece and tried it out. I love the size and no drip mouth.The color is beautiful and its so pretty on my buffet.	Like this pieces I just obtained this pieces and attempts it out. I luv the size and no drip mouths the colorful is beautiful and its however rather on my buffet.

Table 8: Adversarial Samples.

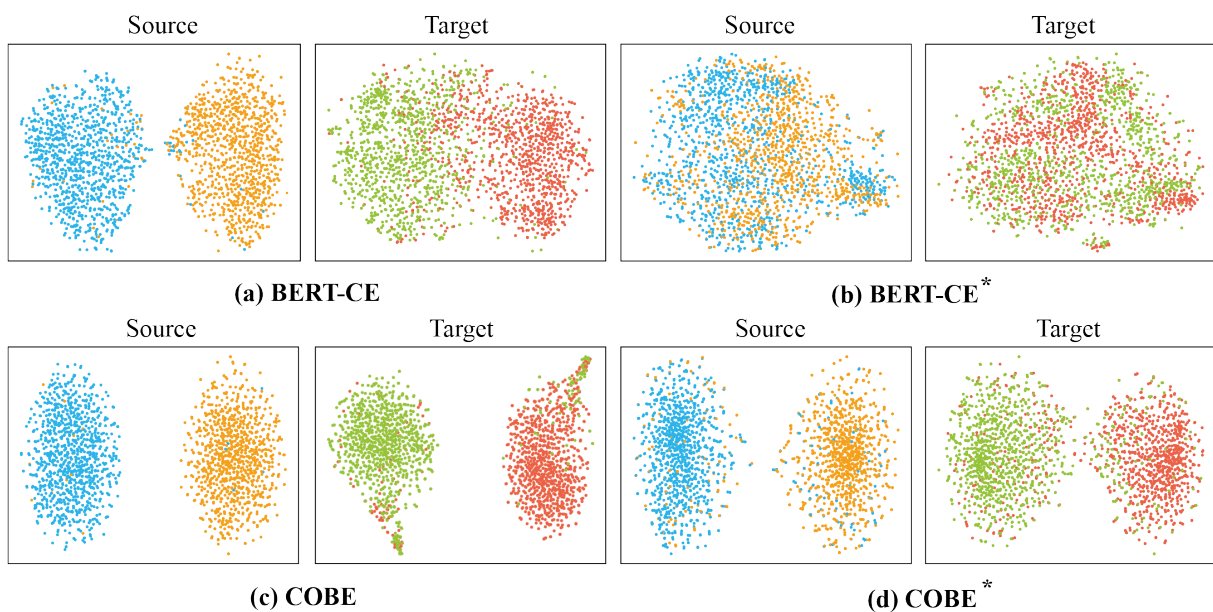


Figure 5: Visualization of sentence representation obtained from BERT and COBE. We use t-SNE to transfer the feature space into two-dimensional space for the B→K task.