

Domain Confused Contrastive Learning for Unsupervised Domain Adaptation

Quanyu Long, Tianze Luo, Wenya Wang and Sinno Jialin Pan

Nanyang Technological University, Singapore

{quanyu001, tianze001, wangwy, sinnopan}@ntu.edu.sg

Abstract

In this work, we study Unsupervised Domain Adaptation (UDA) in a challenging self-supervised approach. One of the difficulties is how to learn task discrimination in the absence of target labels. Unlike previous literature which directly aligns cross-domain distributions or leverages reverse gradient, we propose Domain Confused Contrastive Learning (DCCL) to bridge the source and the target domains via domain puzzles, and retain discriminative representations after adaptation. Technically, DCCL searches for a most domain-challenging direction and exquisitely crafts domain confused augmentations as positive pairs, then it contrastively encourages the model to pull representations towards the other domain, thus learning more stable and effective domain invariances. We also investigate whether contrastive learning necessarily helps with UDA when performing other data augmentations. Extensive experiments demonstrate that DCCL significantly outperforms baselines.

1 Introduction

Pre-trained language models (Devlin et al., 2019; Liu et al., 2019; Yang et al., 2019) have yielded considerable improvements with datasets drawing from various sources. However, the lack of portability of language model to adapt to a new textual domain remains a central issue (Gururangan et al., 2020), especially when the training set and testing set do not follow the same underlying distribution (changing of topic and genres) - usually referred to as domain shift. In this paper, we focus on studying Unsupervised Domain Adaptation (UDA). UDA aims at designing adaptation algorithms that attempt to generalize well on the target domain by learning from both labeled samples from the source domain and unlabeled samples from the target domain. Studying UDA fits real-world scenarios since labeled data in the target domain is usually absent. Moreover, advances in UDA will also help out-of-

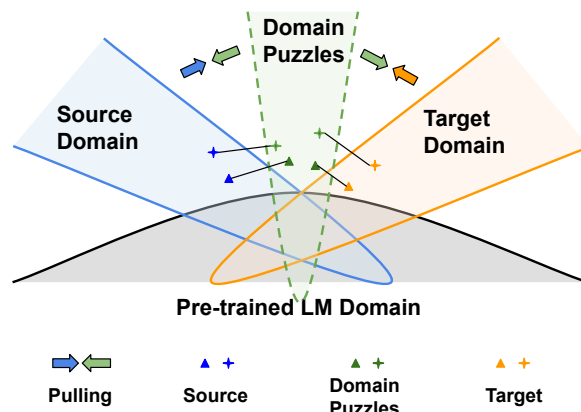


Figure 1: Domain puzzles that are domain-confused and overlook domain-related information, could be regarded as lying in an intermediate domain that aims to pull source and target samples closer to each other and bridge the two domains by learning domain invariant representations.

distribution generalizations (Ramponi and Plank, 2020; Krueger et al., 2021).

Extensive algorithms have been proposed to mitigate the domain shift problem, for example, domain adversarial neural network (DANN) (Ganin et al., 2016) and distribution matching (Zhuang et al., 2015). For DANN, the training process of joint optimization is unstable, requiring extensive effort to tune the hyperparameters (Shah et al., 2018; Du et al., 2020; Karouzos et al., 2021). As for distribution matching, it is very difficult to preserve the discriminative power of the model on the target task while trying to perform instance level alignment (Saito et al., 2017; Lee et al., 2019). To this end, it is essential to develop stable and effective solutions to learn domain invariance and instance-wise matching for UDA.

Recent advances in self-supervised learning (SSL), such as contrastive learning (CL), have been proven effective at instance level by leveraging raw data to define surrogate tasks that help learn repre-

sentations (Chen et al., 2020a; Khosla et al., 2020; He et al., 2020; Chen et al., 2020b). CL benefits from treating instances as classes and conducting data augmentations to generate positive instance pairs. Regarding CL for UDA, previous works mark cross-domain images with the same labels (for example, real and cartoon dogs) as the positive pairs in contrastive loss (Wang et al., 2021; Park et al., 2020). However, such methods are not applicable to NLP tasks because of the massive semantic and syntactic shifts between two cross-domain sentences. Besides, from the domain adaptation perspective, constructing cross-domain positive samples and aligning domain-agnostic pairs have received less emphasis in related literature, since previous works focus on designing label preserving text transformations, such as back-translation, synonym, dropout and their combinations (Qu et al., 2021; Gao et al., 2021).

Confronting with limitations mentioned above, we propose the concept of domain puzzles which discard domain-related information to confuse the model, making it difficult to differentiate which domain these puzzles belong to. Instead of directly seeking matched sentences across the source and target domains which is infeasible, we propose to pull the source (target) data and its corresponding domain puzzles closer to reduce the domain discrepancy, as shown in Fig. 1. A simple idea to craft domain puzzles is to mask domain-specific tokens. However, token-level operations are too discrete and non-flexible to reflect the complex semantic change of natural languages. Hence, we aim to seek better domain puzzles that retain high-confidence predictions and task-discriminative power in the representation space for each training instance.

In this paper, we propose Domain Confused Contrastive Learning (DCCL) to encourage the model to learn similar representations for the original sentence and its curated domain-confused version with contrastive loss. More specifically, we synthesize these domain puzzles by utilizing adversarial examples (Zhu et al., 2020; Jiang et al., 2020). The algorithm will search for an extreme direction that roughly points to the opposite domain and produces most domain-challenging puzzles. We encourage the model to encode original and domain-confused samples closer, gradually pulling examples to the domain decision boundary as training progresses via CL, thus learning the domain invariance. Furthermore, in order to investigate whether CL neces-

sarily benefits UDA, we conduct experiments and find that constructing domain puzzles as paired positive samples is favorable for UDA, however, other data augmentation methods such as back translation (Sennrich et al., 2016; Edunov et al., 2018) do not have the same effect. The experiment results show that the proposed DCCL significantly outperforms all the baselines. We also conduct quantitative experiments to measure the domain discrepancy after adaptation demonstrating that DCCL can decrease the divergence between domains in a self-supervised way. Overall, the paper makes the following contributions:

- First, a new concept of domain puzzles is put forward. We propose to craft the domain puzzles via domain-confused adversarial attack;
- Second, we propose DCCL, which is able to pull source and target samples closer to the crafted domain puzzles. The DCCL is capable of reducing domain shift and learning domain invariance;
- Third, experiments demonstrate that the proposed DCCL surpasses baselines with a large margin. We also conduct analyzing experiments to verify the effectiveness.

2 Preliminaries

2.1 Unsupervised Domain Adaptation

Problem Setup Suppose we have access to a source dataset with n labeled data points $\mathcal{D}^S = \{x_i, y_i\}_{1, \dots, n}$ sampled i.i.d. from the source domain, and a target dataset with m unlabeled points $\mathcal{D}^T = \{x_j\}_{1, \dots, m}$ sampled i.i.d. from the target domain, where x_i, x_j are sequences of tokens, y_i is the class label for x_i . For in-domain training with labeled training instances, the model aims to learn a function $f(x; \theta_f, \theta_y) : x \rightarrow C$. θ_f is the parameter of the deep neural network encoder (e.g., pretrained language model), θ_y denotes parameters that compute the network’s class label predictions, and C is the label set. The model is learned with the following objective:

$$\min_{\theta_f, \theta_y} \sum_{(x, y) \sim \mathcal{D}^S} [\mathcal{L}(f(x; \theta_f, \theta_y), y)]. \quad (1)$$

However, for Unsupervised Domain Adaptation (UDA), the goal of the adaptation algorithm is to learn a discriminative classifier from the source




Figure 2: Overview of the framework of the proposed method DCCL.

domain, which at the same time could generalize well on the target domain by leveraging unlabeled target data and learning a mapping between source and target domains. It is generally acknowledged that the discrepancy between two datasets (domain shift) can be reduced by aligning two distributions (Ben-David et al., 2006; Ben-David et al., 2010). The methods that learn domain invariant features for domain alignment include KL divergence (Zhuang et al., 2015), Maximum Mean Discrepancy (MMD) (Gretton et al., 2012), and Domain Adversarial Neural Network (DANN) (Ganin et al., 2016) (details of DANN can be found in Appendix B). DANN suffers from a vanishing gradient problem (Shen et al., 2018), and the training process is unstable (Shah et al., 2018; Du et al., 2020). Hence, more efficient and stable algorithms are essential for UDA (Wu et al., 2019b).

2.2 Adversarial Training

Adversarial training with perturbations has been shown to significantly improve the performance of the state-of-the-art language models for many natural language understanding tasks (Madry et al., 2018; Zhu et al., 2020; Jiang et al., 2020; Pereira et al., 2021). The algorithm generally considers adversarial attacks with perturbations to word embeddings and minimizes the resultant adversarial loss around the input samples. In a single domain, adversarial training (Goodfellow et al., 2015) is an inner max, outer min adversarial problem with the objective:

$$\min_{\theta_f, \theta_y} \sum_{(x, y) \sim \mathcal{D}} [\max_{\delta} \mathcal{L}(f(x + \delta; \theta_f, \theta_y), y)]. \quad (2)$$

With (2), the standard adversarial training can also be regularized using virtual adversarial training (Miyato et al., 2018), which encourages smoothness in the embedding space. The α_{adv} controls the trade-off between the two losses, usually set to

be 1.

$$\min_{\theta_f, \theta_y} \sum_{(x, y) \sim \mathcal{D}} [\mathcal{L}(f(x; \theta_f, \theta_y), y) + \alpha_{adv} \max_{\delta} \mathcal{L}(f(x + \delta; \theta_f, \theta_y), f(x; \theta_f, \theta_y))]. \quad (3)$$

For (2)(3), the inner maximization can be solved by Projected Gradient Decent (PGD) (Madry et al., 2018) with an additional assumption that the loss function is locally linear. A following iteration can approximate the adversarial perturbation δ :

$$\delta_{t+1} = \Pi_{\|\delta\|_F \leq \epsilon} (\delta_t + \eta \frac{g_y^{adv}(\delta_t)}{\|g_y^{adv}(\delta_t)\|_F}), \quad (4)$$

$$g_y^{adv}(\delta_t) = \nabla_{\delta} \mathcal{L}(f(x + \delta_t; \theta_f, \theta_y), y), \quad (5)$$

where $\Pi_{\|\delta\|_F \leq \epsilon}$ performs a projection onto the ϵ -ball. The advantages of PGD lie in that it only relies on the model itself to produce diverse adversarial samples, enabling the model to generalize better to unseen data.

3 Method

In this section, we focus our discussions on the proposed Domain Confused Contrastive Learning (DCCL) under a sentiment classification scenario. The overall framework of our method is illustrated in Fig. 2. The model will take source labeled and target unlabeled sentences as input. It will then augment the input data with domain puzzles by fabricating adversarial perturbations. With the augmented data, the next step produces a hidden representation for each instance with an encoder which will be further used to produce three losses to train the entire model, namely sentiment classification loss, contrastive loss and consistency loss.

3.1 Crafting domain puzzles

For UDA, Saito et al. (2017) mentions that simply matching the distributions cannot ensure high accuracy on the target domain without the target labels.




Figure 3: Two sentences sampled from Book and Music reviews. Alternatively, we can match original sentences with its degraded masked versions.


Moreover, it may cause negative transfer, deteriorating knowledge transfer from source domain to the target domain (Wang et al., 2019). Even if the matched sentences have the same label, due to huge syntactic and semantic shift, instance-based matching strategies that align examples from different domains will introduce noises for pre-trained language models, for example, aligning source domain and target domain sentences in Fig. 3.

Alternatively, we can locate and mask domain-specific tokens which are related to sentence topics and genres. Since sentences in the green box of Fig. 3 become domain-agnostic, we refer to those domain-confused sentences (one cannot tell which domain these sentences belong to) as domain puzzles. Matching distributions between the source domain and the domain puzzles, as well as the target domain and the domain puzzles, will also make language models produce domain invariant representations.

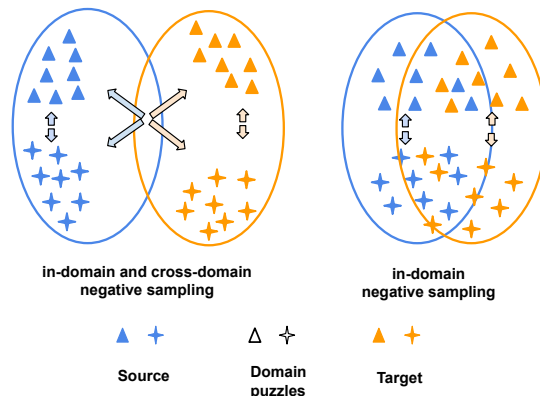
However, the domain-specific tokens are not always evident, due to the discrete nature of natural languages, it is challenging to decide correct tokens to mask without hurting the semantics especially when the sentences are complicated¹. Hence, we seek domain puzzles in the representation space and introduce adversarial perturbations, because we can rely on the model itself to produce diverse but targeted domain puzzles. Note that the purpose of adversarial attack here is not to enhance the robustness, but to construct exquisitely produced perturbations for a better domain invariance in the representation space.

To generate domain-confused augmentations, we adopt adversarial attack with perturbations for domain classification. The loss for learning a domain classifier with adversarial attack can be speci-

¹Masking is not our focus in this paper. More detailed implementation can be found in section 4.4. We will investigate how to extract better domain tokens in our future work.



(a) Positive sampling



(b) Negative sampling

Figure 4: Proposed adaptation sampling method

fied as follows:

$$\mathcal{L}_{\text{domain}} = \mathcal{L}(f(x; \theta_f, \theta_d), d) + \alpha_{adv} \mathcal{L}(f(x + \delta; \theta_f, \theta_d), f(x; \theta_f, \theta_d)), \quad (6)$$

$$\delta = \Pi_{\|\delta\|_F \leq \epsilon} \left(\delta_0 + \eta \frac{g_d^{adv}(\delta_0)}{\|g_d^{adv}(\delta_0)\|_F} \right), \quad (7)$$

where δ_0 is the initialized noise, θ_d is the parameter corresponding to the computation of the domain classification, and d is the domain label. Due to additional overhead incurred during fine-tuning large pre-trained language models, the number of iterations for perturbation estimation is usually 1 (Jiang et al., 2020; Pereira et al., 2021), as shown in Eq. 7. We synthesize the perturbation δ by searching for an extreme direction that perplexes the domain classifier most in the embedding space, and $f(x + \delta; \theta_f)$ is the crafted domain puzzles encoded by the language model.

3.2 Learning invariance with domain puzzles

After acquiring domain puzzles, simply applying distribution matching will sacrifice discriminative knowledge learned from the source domain (Saito et al., 2017; Lee et al., 2019), and instance-based matching will also overlook global intra-domain information. To learn sentiment-wise discriminative

representations in the absence of the target labels, we propose to learn domain invariance via contrastive learning (CL). In general, CL benefits from the definition of the augmented positive and negative pairs by treating instances as classes (Chen et al., 2020a; Khosla et al., 2020; He et al., 2020; Chen et al., 2020b). Furthermore, the contrastive loss encourages the positive pairs to be close to each other and negative pairs to be far apart. Specifically, maximizing the similarities between positive pairs learns an invariant instance-based representation, and minimizing the similarities between negative pairs learns a uniformly distributed representation from a global view, making instances gathered near the task decision boundary away from each other (Saunshi et al., 2019; Grill et al., 2020). This will help to enhance task discrimination of the learned model.

For positive pairs, intuitively, we hope that the model could encode the original sentence and most domain-challenging examples to be closer in the representation space, gradually pulling examples to the domain decision boundary as training progresses. For negative sampling, it widens the sentiment decision boundary and promotes better sentiment-wise discriminative features for both domains. However, for cross-domain negative sampling, the contrastive loss may push the negative samples in the target (source) domain away from the anchor in the source (target) domain (see Fig. 4 (b) left). This is contradictory to the objective of domain puzzles which try to pull different domains closer. To avoid the detriment of cross-domain repulsion, excluding samples with different domains from the negative set is of great importance. Therefore, we write the following contrastive infoNCE loss (Chen et al., 2020a) as follow:

$$\mathcal{L}_{\text{contrast}} = -\frac{1}{N} \sum_i \log \frac{\exp(s(z_i, z'_i)/\tau)}{\sum_k \mathbb{1}_{k \neq i} \exp(s(z_i, z_k)/\tau)}, \quad (8)$$

where N is the mini batch size with samples from the same domain, $z_i = g(f(x_i; \theta_f))$, and $g(\cdot)$ is one hidden layer projection head. We denote $x' = x + \delta$ as the domain puzzle augmentation, $s(\cdot)$ computes cosine similarity, $\mathbb{1}_{k \neq i}$ is the indicator function, and τ is the temperature hyperparameter.

3.3 Consistency Regularization

Given perturbed embedding $x + \delta$, which is crafted based on domain classification, we also encourage the model to produce consistent sentiment predictions with that of the original instance $f(x; \theta_f, \theta_y)$.

Algorithm 1 DCCL

Input: For simplicity, θ is the parameter of the whole model. T : the total number of iterations, $(x, y) \sim \mathcal{D}^S$: source dataset with sentiment label y , $(x, d) \sim \mathcal{D}^S \mathcal{D}^D$: source and target dataset with domain label d , K : the number of iterations for updating δ , σ^2 : the initialized variance, ϵ : perturbation bound, η : the step size, γ : global learning rate, N : batch size, τ : temperature, $g(\cdot)$: one hidden layer projection head. α_{adv} , α , λ and β : weighting factor.

```

1: for epoch = 1, ..., T do
2:   for minibatch  $N$  do
3:      $\delta \leftarrow \mathcal{N}(0, \sigma^2 I)$ 
4:     for  $m = 1, \dots, K$  do
5:        $g_d^{adv} \leftarrow \nabla_{\delta} \mathcal{L}(f(x + \delta; \theta), d)$ 
6:        $\delta \leftarrow \Pi_{\|\delta\|_F \leq \epsilon}(\delta + \eta g_d^{adv} / \|g_d^{adv}\|_F)$ 
7:     end for
8:      $\mathcal{L}_{\text{domain}} \leftarrow \mathcal{L}(f(x; \theta), d)$ 
9:        $+ \alpha_{adv} \mathcal{L}(f(x + \delta; \theta), d)$ 
10:     $z = g(f(x; \theta))$ 
11:     $z' = g(f(x + \delta; \theta))$ 
12:    for  $i = 1, \dots, N$  and  $j = 1, \dots, N$  do
13:       $s'_i = z_i^\top z'_j / \|z_i\| \|z'_j\|$ 
14:       $s_{i,j} = z_i^\top z_j / \|z_i\| \|z_j\|$ 
15:    end for
16:     $\mathcal{L}_{\text{contrast}} \leftarrow -\frac{1}{N} \sum_i \log \frac{\exp(s'_i/\tau)}{\sum_j \mathbb{1}_{j \neq i} \exp(s_{i,j}/\tau)}$ 
17:     $\mathcal{L}_{\text{consist}} \leftarrow \mathcal{L}(f(x; \theta), f(x + \delta; \theta))$ 
18:     $g_{\theta} \leftarrow \nabla_{\theta} \mathcal{L}(f(x; \theta), y) + \alpha \nabla_{\theta} \mathcal{L}_{\text{domain}}$ 
19:       $+ \lambda \nabla_{\theta} \mathcal{L}_{\text{contrast}} + \beta \nabla_{\theta} \mathcal{L}_{\text{consist}}$ 
20:     $\theta \leftarrow \theta - \gamma g_{\theta}$ 
21:  end for
22: end for
Output:  $\theta$ 

```

For this, we minimize the symmetric KL divergence, which is formulated as:

$$\mathcal{L}_{\text{consist}} = \mathcal{L}(f(x; \theta_f, \theta_y), f(x + \delta; \theta_f, \theta_y)). \quad (9)$$

For overall training objective, we train the neural network in an end-to-end manner with a weighted sum of losses as follows.

$$\min_{\theta_f, \theta_y, \theta_d} \sum_{(x,y) \sim \mathcal{D}^S} \mathcal{L}(f(x; \theta_f, \theta_y), y) + \sum_{(x,y) \sim \mathcal{D}^S, \mathcal{D}^T} [\alpha \mathcal{L}_{\text{domain}} + \lambda \mathcal{L}_{\text{contrast}} + \beta \mathcal{L}_{\text{consist}}]. \quad (10)$$

Details of proposed DCCL are summarized in Algorithm 1.

4 Experiments

4.1 Datasets

Amazon Benchmark (Blitzer et al., 2007)². We conduct experiments on this dataset for completeness since most of previous works report results on it. The dataset contains four domains: Book (BK), DVD (D), Electronics (E) and Kitchen housewares (K). There are 2,000 balanced labeled data for each domain, we randomly select 20,000 unlabeled reviews for BK, D and E. For K, only 17,856 unlabeled reviews are available.

Amazon review dataset (He et al., 2018)³. This dataset considers neutral instances which may not bias the dataset and bring more challenges⁴. The dataset also contains four domains: Book (BK), Electronics (E), Beauty (BT), and Music (M). Following He et al. (2018), we treat set 1 as labeled dataset containing 6,000 instances, and treat set 2 as unlabeled dataset which also contains 6,000 instances. More details about two datasets can be found in Appendix A.

4.2 Experiment Settings

For unsupervised adaptation setting, we should not have access to target labeled data at the training phase, so trained models with minimum classification error on the source validation set is saved for evaluation. At this point, we suppose a good model that generalizes well on the target domain is able to reach high performance on both validation and test set at the same time. We evaluate our model with 5 runs in all experiments, and we report the average score, standard deviation and paired t-test results.

4.3 Implementation Details

For pre-trained language model, we use BERT base uncased (Devlin et al., 2019) as the basis for all experiments. The max length is set to 512. For optimizer, we use AdamW (Kingma and Ba, 2015; Loshchilov and Hutter, 2019) with weight decay 0.01 (for BERT baseline, we set 1e-4). We set the learning rate as 1e-5, and we use a linear scheduler with warm-up steps 0.1 of total training steps.

We set the number of adversarial iterations to be 1, adversarial weighting factor $\alpha_{adv} = 1$ and we use l_2 norm to compute projections. We also

follow Zhu et al. (2020) to set other adversarial hyperparameters, e.g., adversarial step size $\eta = 5e-2$, perturbation bound $\epsilon = 5e-2$. For weighting factors, we set $\alpha = 1e-3$, $\lambda = 3e-2$ and $\beta = 5$. We train the model with 8 epochs, temperature τ and batch size N will be discussed later. Each adaptation requires half one hour on one A-100.

4.4 Baselines

BERT base: Fine-tune BERT on the source, without using target unlabeled data, then directly evaluate on the target labeled data. **KL**: Use symmetric KL-divergence loss of embedded instances between the source and target domains (Zhuang et al., 2015). **MMD**: Maximum Mean Discrepancy loss (Gretton et al., 2012) measures the distance based on the notion of embedding probabilities in a reproducing kernel Hilbert space. We implement a gaussian kernel which is a common choice. **DANN**: The adaptation rate is $\lambda = \frac{2}{1+\exp(-\gamma p)} - 1$, $p = \frac{t}{T}$, where t and T are the number of current training steps and total steps. γ requires careful tuning within [0.05, 0.1, 0.15, 0.2]. **backtrans+CL**: To investigate the effectiveness of domain puzzles, we implement Back translation (Sennrich et al., 2016; Edunov et al., 2018), which is one of the widely used data augmentation techniques. We utilize en-de translation model pre-trained on WMT19 and released in fairseq (Ott et al., 2019). The model is trained with contrastive loss on the source and target domain respectively. **mask+CL**: We mask domain specific tokens to make the augmentation become domain-agnostic. Since information extraction is not our focus, we identify domain specific tokens via a simple frequency-ratio method (Li et al., 2018; Wu et al., 2019a): $s(u, d) = \frac{\text{count}(u, \mathcal{D}^d) + \lambda}{\sum_{d' \in \mathcal{D}, d' \neq d} \text{count}(u, \mathcal{D}^{d'}) + \lambda}$, where $\text{count}(u, \mathcal{D}^d)$ represents the number of times a token u appears in domain d . Smoothing λ is set to 1. When $s(u, d)$ is larger than 5, we mark a token u as a domain specific token. Through counting, the number of masked tokens accounted for 0.06 of the total length. **mask**: To investigate the effectiveness of contrastive loss, after masking domain-specific tokens, we further let the model train with augmented data without contrastive loss. **R-PERL** (Ben-David et al., 2020) is a pivot-based method, and **DAAT** (Du et al., 2020) combines DANN and post-training. All the methods in Tabel 1 and Table 2 are implemented based on BERT model for fair comparisons.

²<https://www.cs.jhu.edu/~mdredze/datasets/sentiment/>

³<https://github.com/ruidan/DAS>

⁴The original crawled reviews contain star ratings (1 to 5 stars). Comparing with Amazon Benchmark which discards the neutral class, this dataset labels them with rating < 3 , > 3 , $= 3$ as negative, positive, and neutral respectively.

Method	E→BK	BT→BK	M→BK	BK→E	BT→E	M→E	BK→BT	E→BT	M→BT	BK→M	E→M	BT→M	Ave.
BERT base	64.34 [±] _{0.3}	65.87 [±] _{1.9}	64.12 [±] _{1.7}	52.25 [±] _{1.8}	66.01 [±] _{1.0}	59.49 [±] _{1.1}	55.01 [±] _{1.7}	66.02 [±] _{1.3}	56.50 [±] _{2.1}	55.66 [±] _{1.5}	61.34 [±] _{0.8}	62.34 [±] _{1.2}	60.74 _{1.5}
KL	64.77 [±] _{0.8}	67.03 [±] _{0.7}	65.13 [±] _{1.0}	56.96 [±] _{1.2}	65.43 [±] _{1.5}	60.30 [±] _{0.7}	58.47 [±] _{0.9}	66.22 [±] _{0.5}	58.80 [±] _{1.3}	57.53 [±] _{0.7}	60.02 [±] _{1.4}	64.41 [±] _{1.1}	62.09 _{1.0}
MMD	65.41 [±] _{0.7}	68.54 [±] _{1.1}	64.77 [±] _{1.3}	58.84 [±] _{1.6}	66.39 [±] _{1.5}	59.29 [±] _{1.9}	60.75 [±] _{1.5}	66.50 [±] _{0.9}	57.42 [±] _{1.3}	59.37 [±] _{1.4}	59.47 [±] _{1.0}	65.03 [±] _{1.6}	62.65 _{1.3}
DANN	66.43 [±] _{2.4}	67.74 [±] _{2.9}	65.64 [±] _{3.7}	54.31 [±] _{4.1}	65.93 [±] _{3.0}	59.91 [±] _{1.9}	58.11 [±] _{3.5}	67.19 [±] _{0.5}	56.00 [±] _{2.7}	54.02 [±] _{2.9}	60.36 [±] _{3.1}	63.09 [±] _{2.7}	61.56 _{2.8}
back-trans+CL	67.16 [±] _{1.4}	68.31 [±] _{0.8}	67.78 [±] _{1.8}	59.55 [±] _{1.1}	66.51 [±] _{0.3}	60.13 [±] _{1.4}	60.92 [±] _{0.7}	68.54 [±] _{0.8}	60.44 [±] _{0.9}	60.21 [±] _{1.1}	61.73 [±] _{1.3}	64.29 [±] _{0.7}	63.79 _{1.0}
mask	65.17 [±] _{1.1}	66.91 [±] _{1.4}	65.59 [±] _{0.9}	52.74 [±] _{1.5}	67.14 [±] _{1.9}	60.48 [±] _{0.7}	54.76 [±] _{1.6}	66.50 [±] _{1.4}	56.61 [±] _{1.7}	55.64 [±] _{1.6}	62.17 [±] _{0.9}	63.09 [±] _{1.0}	61.40 _{1.3}
mask+CL	68.80 [±] _{1.2}	69.93 [±] _{0.7}	70.65 [±] _{0.9}	60.50 [±] _{1.3}	68.02 [±] _{0.7}	60.80 [±] _{1.2}	62.17 [±] _{0.8}	69.17 _{0.5}	62.06 [±] _{1.5}	60.56 [±] _{1.4}	63.94 _{0.7}	65.40 _{0.6}	65.17 _{0.9}
DCCL	70.33 _{0.3}	70.92 _{0.6}	71.11 _{0.7}	62.36 _{0.7}	68.41 _{0.2}	62.11 _{0.6}	64.13 _{1.4}	69.33 _{0.4}	65.40 _{0.8}	64.67 _{1.7}	64.67 _{1.0}	66.70 _{1.4}	66.68 _{0.8}

Table 1: Accuracy (%) results of Amazon review dataset. For example, E→BK denotes training on Electronics (E) and adapting to Book (BK). All the results are reported with average and standard deviation in 5 runs. * indicates the DCCL improvements are significant with $p < 0.05$. There are four domains available in the dataset, therefore we have 12 adaptation groups of tasks.

Method	D→BK	K→BK	E→BK	BK→D	K→D	E→D	BK→K	D→K	E→K	BK→E	D→E	K→E	Ave.
R-PERL	85.6	83.0	83.9	87.8	85.6	84.8	90.2	90.4	91.2	87.2	89.3	91.2	87.5
DAAT	90.86	87.98	88.91	89.70	88.81	90.13	90.75	90.50	93.18	89.57	89.30	91.72	90.12
BERT base	89.76 [±] _{0.3}	88.31 [±] _{0.4}	88.09 [±] _{0.9}	89.27 [±] _{0.5}	87.68 [±] _{0.8}	88.41 [±] _{1.2}	88.27 [±] _{0.4}	87.67 [±] _{0.5}	92.16 _{0.7}	87.18 [±] _{0.4}	86.91 [±] _{0.7}	91.22 _{0.8}	88.74 _{0.6}
KL	89.23 [±] _{0.5}	87.61 [±] _{0.5}	88.37 [±] _{0.7}	89.49 [±] _{0.6}	88.03 [±] _{0.5}	88.56 [±] _{0.7}	88.77 [±] _{0.5}	87.89 [±] _{0.7}	91.21 _{1.3}	88.62 [±] _{0.4}	87.03 [±] _{0.9}	90.34 [±] _{1.5}	88.76 _{0.7}
MMD	88.68 [±] _{0.7}	88.61 _{0.3}	88.27 [±] _{1.0}	88.38 [±] _{1.0}	89.09 _{0.5}	89.31 [±] _{0.6}	87.74 [±] _{0.8}	89.04 [±] _{0.8}	91.20 _{1.0}	89.49 [±] _{0.3}	88.70 [±] _{0.5}	89.79 [±] _{1.2}	89.03 _{0.6}
DANN	88.64 [±] _{1.3}	86.09 [±] _{1.7}	87.91 [±] _{1.5}	89.56 _{1.0}	89.11 _{0.9}	89.19 [±] _{1.1}	89.68 [±] _{1.3}	88.73 [±] _{1.8}	91.41 [±] _{1.3}	87.30 [±] _{1.6}	88.53 [±] _{0.6}	90.81 [±] _{1.1}	88.91 _{1.3}
mask+CL	90.74 _{0.3}	88.28 _{0.7}	89.44 [±] _{0.4}	90.31 _{0.2}	88.74 [±] _{0.5}	89.72 [±] _{0.8}	90.40 [±] _{0.4}	89.57 [±] _{0.4}	92.49 _{0.4}	89.91 [±] _{0.3}	89.65 _{0.4}	91.19 _{0.6}	90.08 _{0.5}
DCCL	91.17 _{0.3}	88.53 _{0.4}	89.70 _{0.5}	90.03 _{0.4}	89.53 _{0.3}	90.49 _{0.5}	91.05 _{0.5}	90.78 _{0.4}	92.54 _{0.5}	90.42 _{0.3}	89.55 _{0.4}	91.93 _{0.5}	90.48 _{0.4}

Table 2: Accuracy (%) results of Amazon Benchmark. Results of R-PERL and DAAT are taken from Ben-David et al. (2020) and Du et al. (2020) respectively.

4.5 Results

Will contrastive learning necessarily help the Unsupervised Domain Adaptation?

As discussed earlier, when performing contrastive learning on source labeled examples and target unlabeled examples respectively, it learns a uniformly distributed representation and helps promote better discriminative features for both domains. However, for some adaptation tasks in Table 1, for example, BT→E, M→E, and E→M, back-trans+CL shows that contrastive learning only gains marginal benefit. When masking domain-specific tokens and pulling original sentence representations to those domain puzzles, the effect of contrastive learning becomes more apparent (mask+CL with average score 65.17 compared with back-trans+CL 63.79 and mask 61.40 in Table.1). This finding helps to explain that choices of positive examples are critical and domain confused augmentations will further benefit adaptation.

DCCL outperforms baselines.

From Table 1 we can observe that the proposed DCCL outperforms all other methods with a large margin and $p < 0.05$ using paired t-test, and 5.94% improvement over BERT base. From Table 2, we can also observe 1.74% improvement over BERT base, and DCCL also surpasses state-of-the-art

methods R-PERL (Ben-David et al., 2020) and DAAT (Du et al., 2020). We note that Amazon Benchmark dataset is quite easy, since it discards neutral instances, BERT base model has already achieved high scores on this dataset. Besides, we observe that the effect of distribution matching methods (KL and MMD) is limited on two datasets. The reason might be that pre-trained language models trained with massive and heterogeneous corpora already have strong generalization ability. Learning such cross-domain and instance-based matching will bring perplexity to language models and sacrifice task discrimination (Saito et al., 2017; Lee et al., 2019). On the contrary, the proposed DCCL retains such information in a self-supervised way. Furthermore, we notice that DANN is very unstable, besides adaptation rate λ , the model is also sensitive to other hyperparameters such as learning rate and training epochs because the performance on the target domain will keep decreasing when training with longer steps. Hence, it is difficult for the model to achieve the lowest error rates on both the source and target domains simultaneously. Compared to DANN, DCCL is much more stable and has lower standard deviations on most adaptation tasks.

Contrastive learning designs.

We explore different hyperparameter values for the

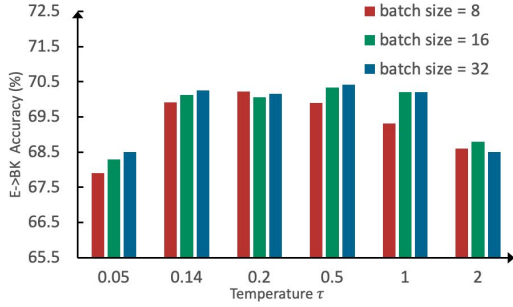


Figure 5: Hyperparameters for contrastive learning

proposed DCCL in E→BK, as is shown in Fig. 5. We find that a temperature of 0.5 combined with the batch size 32 can achieve the best performance. We also notice that setting the temperature too high or too low can significantly affect adaptation performances, while a larger batch size usually brings a relatively smaller improvement.

4.6 Ablation Studies

We conduct an ablation study on each component in Eq. 10 to inspect the contribution of contrastive learning, as is shown in Table 3. We can see that every single component can render a better adaptation performance. In particular, the effectiveness of $\mathcal{L}_{\text{contrast}}$ is observable with 3~5 performance gain compared to baselines (row 4 vs. row 2, and row 5 vs. row 3). When training curated domain puzzles as simple augmentations without contrastive loss, we can observe only a slight improvement (row 2 vs. row 1). This result demonstrates that the performance gain brought by DCCL does not come from data augmentation, and learning robustness against adversarial attacks will not largely help adaptation.

$\mathcal{L}_{\text{domain}}$	$\mathcal{L}_{\text{consist}}$	$\mathcal{L}_{\text{contrast}}$	E→BK	M→BT
			64.65	55.85
✓			67.23	58.9
✓	✓		67.85	59.1
✓		✓	70.12	64.73
✓	✓	✓	70.21	64.87

Table 3: Ablation studies of DCCL on each component

5 Analysis

5.1 Visualization

We perform visualizations for trained representations as illustrated in Fig. 6. When training with the source domain and then adapting to the target domain (BERT-base), we can observe a considerable

domain shift for BERT encoder on this amazon review dataset. Moreover, as mentioned before, continuing training DANN with larger epochs will substantially drop the score (from the highest point (DANN-best) to the lowest point (DANN-worst)). However, we can also see that DCCL mitigates domain shift but remains good sentiment discrimination on the target domain.

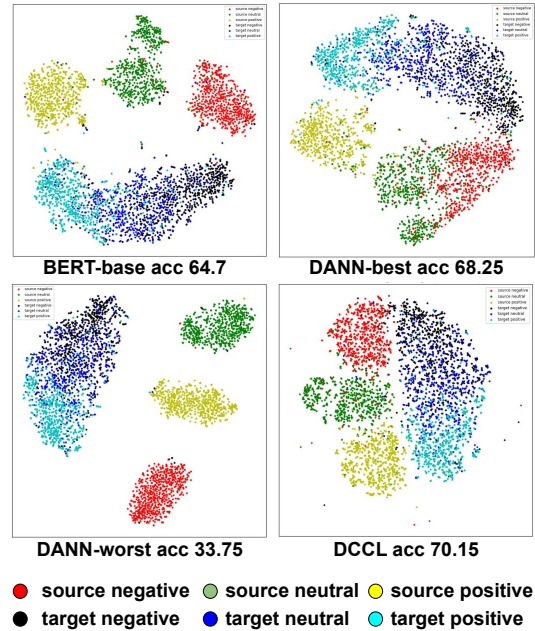


Figure 6: t-SNE visualization for E→BK task

5.2 Quantitative Results

\mathcal{A} -distance measures domain discrepancies (Ben-David et al., 2006; Ben-David et al., 2010), with the definition as $d_{\mathcal{A}} = 2(1 - 2\epsilon)$, where ϵ is the domain classification error. To fairly compare with \mathcal{A} -distance of the baselines, we use linear SVM to calculate ϵ following previous work (Saito et al., 2017; Du et al., 2020). We randomly select 2,000 instances for both source and target domain and split them with 1:1 for train and test for the SVM model. From Fig. 7, we can observe that DCCL can learn a good balance between sentiment classification and domain discrepancy, compared to DANN-best and DANN-worst.

6 Related Work

Unsupervised Domain Adaptation

UDA in NLP has the following approaches: (1) Pivot-based methods use unlabeled data from both domains, trying to discover characteristics that are similar (Pan et al., 2010). Some recent works




Figure 7: \mathcal{A} -distance for E→BK and M→BT tasks

extend pivots with autoencoders and contextualized embeddings (Ziser and Reichart, 2017; Miller, 2019; Ben-David et al., 2020). (2) Pseudo-labeling leverages a trained classifier to predict labels on unlabeled examples, which are subsequently considered as gold labels (Yarowsky, 1995; Zhou and Li, 2005; McClosky et al., 2006). Recent works also combine this technique with pre-trained language models (Lim et al., 2020; Ye et al., 2020). (3) Data selection methods adopt domain similarity metrics to find the best match for each data and use curriculum learning for large pre-trained models (Ma et al., 2019; Aharoni and Goldberg, 2020). A recent method adopt distance-bandit (Guo et al., 2020) for similarity metric. (4) Domain Adversarial Neural Networks (Ganin et al., 2016). Some approaches leverage Wasserstein distance to stabilize adversarial training (Shen et al., 2018; Shah et al., 2018), and combine it with post-training which can produce better adversarial results (Du et al., 2020). (5) Adaptive pre-training is a more straightforward but effective method (Gururangan et al., 2020; Han and Eisenstein, 2019; Karouzos et al., 2021) by leveraging the objective of masked language model (MLM). A wide range of pre-training methods for domain adaptation (multi-phase, multi-task) are put forward (Han and Eisenstein, 2019; Karouzos et al., 2021).

Contrastive learning

CL has recently gained popularity as a reliable approach for unsupervised representation learning. It is generally acknowledged that a good representation should distinguish itself from other instances while identifying similar instances. For Computer Vision, there are approaches obtaining augmented images using transformations including cropping, rotation, etc. (Chen et al., 2020a; Khosla et al., 2020; He et al., 2020; Chen et al., 2020b). As for Natural Language Processing, many works study different label-preserving augmentations, such as back-translations, synonyms, adversaries, dropout,

and their combinations (Qu et al., 2021; Gao et al., 2021). In addition, many pre-trained language models trained with contrastive loss are also released. DeCLUTR (Giorgi et al., 2021) and CLEAR (Wu et al., 2020) jointly train the model with a contrastive objective and a masked language model setting. ConSERT (Yan et al., 2021) overcomes the collapse problem of BERT-derived sentence representations and makes them more suitable for downstream applications by using unlabeled texts.

7 Conclusion

In this work, we put forward a new concept, domain puzzles, which can be crafted through domain-specific token mask and domain-confused adversarial attacks. And we provide a more stable and effective solution to learn domain invariance for unsupervised domain adaptation. The proposed method DCCL surpasses baselines with a large margin by mitigating domain shift without losing discriminative power on the target domain. Moreover, the proposed framework can also be extended to other NLP tasks demanding adaptations, and we leave this for future work.

Acknowledgements

This work is supported by the 2020 Microsoft Research Asia collaborative research grant.

References

- Roe Aharoni and Yoav Goldberg. 2020. [Unsupervised domain clusters in pretrained language models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7747–7763, Online. Association for Computational Linguistics.
- Eyal Ben-David, Carmel Rabinovitz, and Roi Reichart. 2020. [PERL: Pivot-based domain adaptation for pre-trained deep contextualized embedding models](#). *Transactions of the Association for Computational Linguistics*, 8:504–521.
- Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. 2010. A theory of learning from different domains. *Machine learning*, 79(1):151–175.
- Shai Ben-David, John Blitzer, Koby Crammer, and Fernando Pereira. 2006. [Analysis of representations for domain adaptation](#). In *Advances in Neural Information Processing Systems 19, Proceedings of the Twentieth Annual Conference on Neural Information Processing Systems, Vancouver, British Columbia, Canada, December 4-7, 2006*, pages 137–144. MIT Press.

- 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, Prague, Czech Republic. Association for Computational Linguistics.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E. Hinton. 2020a. [A simple framework for contrastive learning of visual representations](#). In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 1597–1607. PMLR.
- Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. 2020b. [Improved baselines with momentum contrastive learning](#). *ArXiv preprint*, abs/2003.04297.
- 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, Minneapolis, Minnesota. Association for Computational Linguistics.
- 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, Online. Association for Computational Linguistics.
- Sergey Edunov, Myle Ott, Michael Auli, and David Grangier. 2018. [Understanding back-translation at scale](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 489–500, Brussels, Belgium. Association for Computational Linguistics.
- 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.
- 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, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- John Giorgi, Osvald Nitski, Bo Wang, and Gary Bader. 2021. [DeCLUTR: Deep contrastive learning for unsupervised textual representations](#). 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 879–895, Online. Association for Computational Linguistics.
- Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. [Explaining and harnessing adversarial examples](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Arthur Gretton, Karsten M Borgwardt, Malte J Rasch, Bernhard Schölkopf, and Alexander Smola. 2012. A kernel two-sample test. *The Journal of Machine Learning Research*, 13(1):723–773.
- Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre H. Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Ávila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, Bilal Piot, Koray Kavukcuoglu, Rémi Munos, and Michal Valko. 2020. [Bootstrap your own latent - A new approach to self-supervised learning](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Han Guo, Ramakanth Pasunuru, and Mohit Bansal. 2020. [Multi-source domain adaptation for text classification via distancenet-bandits](#). In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 7830–7838. AAAI Press.
- Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. 2020. [Don’t stop pretraining: Adapt language models to domains and tasks](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8342–8360, Online. Association for Computational Linguistics.
- Xiaochuang Han and Jacob Eisenstein. 2019. [Unsupervised domain adaptation of contextualized embeddings for sequence labeling](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4238–4248, Hong Kong, China. Association for Computational Linguistics.
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross B. Girshick. 2020. [Momentum contrast for unsupervised visual representation learning](#). In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 9726–9735. IEEE.
- Ruidan He, Wee Sun Lee, Hwee Tou Ng, and Daniel Dahlmeier. 2018. [Adaptive semi-supervised learning](#)

- for cross-domain sentiment classification. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3467–3476, Brussels, Belgium. Association for Computational Linguistics.
- Haoming Jiang, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and Tuo Zhao. 2020. **SMART: Robust and efficient fine-tuning for pre-trained natural language models through principled regularized optimization**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2177–2190, Online. Association for Computational Linguistics.
- Constantinos Karouzos, Georgios Paraskevopoulos, and Alexandros Potamianos. 2021. **UDALM: Unsupervised domain adaptation through language modeling**. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2579–2590, Online. Association for Computational Linguistics.
- Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschiot, Ce Liu, and Dilip Krishnan. 2020. **Supervised contrastive learning**. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Diederik P. Kingma and Jimmy Ba. 2015. **Adam: A method for stochastic optimization**. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- David Krueger, Ethan Caballero, Jörn-Henrik Jacobsen, Amy Zhang, Jonathan Binas, Dinghui Zhang, Rémi Le Priol, and Aaron C. Courville. 2021. **Out-of-distribution generalization via risk extrapolation (rex)**. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 5815–5826. PMLR.
- Seungmin Lee, Dongwan Kim, Namil Kim, and Seong-Gyun Jeong. 2019. **Drop to adapt: Learning discriminative features for unsupervised domain adaptation**. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 91–100. IEEE.
- Juncen Li, Robin Jia, He He, and Percy Liang. 2018. **Delete, retrieve, generate: a simple approach to sentiment and style transfer**. 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 1865–1874, New Orleans, Louisiana. Association for Computational Linguistics.
- KyungTae Lim, Jay Yoon Lee, Jaime G. Carbonell, and Thierry Poibeau. 2020. **Semi-supervised learning on meta structure: Multi-task tagging and parsing in low-resource scenarios**. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 8344–8351. AAAI Press.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, M. Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. **Roberta: A robustly optimized bert pretraining approach**. *ArXiv*, abs/1907.11692.
- Ilya Loshchilov and Frank Hutter. 2019. **Decoupled weight decay regularization**. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Xiaofei Ma, Peng Xu, Zhiguo Wang, Ramesh Nallapati, and Bing Xiang. 2019. **Domain adaptation with BERT-based domain classification and data selection**. In *Proceedings of the 2nd Workshop on Deep Learning Approaches for Low-Resource NLP (DeepLo 2019)*, pages 76–83, Hong Kong, China. Association for Computational Linguistics.
- Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. 2018. **Towards deep learning models resistant to adversarial attacks**. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- Julian J. McAuley, Christopher Targett, Qinfeng Shi, and Anton van den Hengel. 2015. **Image-based recommendations on styles and substitutes**. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval, Santiago, Chile, August 9-13, 2015*, pages 43–52. ACM.
- David McClosky, Eugene Charniak, and Mark Johnson. 2006. **Reranking and self-training for parser adaptation**. In *Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics*, pages 337–344, Sydney, Australia. Association for Computational Linguistics.
- Timothy Miller. 2019. **Simplified neural unsupervised domain adaptation**. 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 414–419, Minneapolis, Minnesota. Association for Computational Linguistics.
- Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. 2018. **Virtual adversarial training:**

- a regularization method for supervised and semi-supervised learning. *IEEE transactions on Pattern Analysis and Machine Intelligence*, 41(8):1979–1993.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. [fairseq: A fast, extensible toolkit for sequence modeling](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 48–53, Minneapolis, Minnesota. Association for Computational Linguistics.
- Sinno Jialin Pan, Xiaochuan Ni, Jian-Tao Sun, Qiang Yang, and Zheng Chen. 2010. [Cross-domain sentiment classification via spectral feature alignment](#). In *Proceedings of the 19th International Conference on World Wide Web, WWW 2010, Raleigh, North Carolina, USA, April 26-30, 2010*, pages 751–760. ACM.
- Changhwa Park, Jonghyun Lee, Jaeyoon Yoo, Minhoe Hur, and Sungroh Yoon. 2020. [Joint contrastive learning for unsupervised domain adaptation](#). *ArXiv preprint*, abs/2006.10297.
- Lis Pereira, Xiaodong Liu, Hao Cheng, Hoifung Poon, Jianfeng Gao, and Ichiro Kobayashi. 2021. [Targeted adversarial training for natural language understanding](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5385–5393, Online. Association for Computational Linguistics.
- Yanru Qu, Dinghan Shen, Yelong Shen, Sandra Sajeew, Weizhu Chen, and Jiawei Han. 2021. [Coda: Contrast-enhanced and diversity-promoting data augmentation for natural language understanding](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Alan Ramponi and Barbara Plank. 2020. [Neural unsupervised domain adaptation in NLP—A survey](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6838–6855, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Kuniaki Saito, Yoshitaka Ushiku, and Tatsuya Harada. 2017. [Asymmetric tri-training for unsupervised domain adaptation](#). In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pages 2988–2997. PMLR.
- Nikunj Saunshi, Orestis Plevrakis, Sanjeev Arora, Mikhail Khodak, and Hrishikesh Khandeparkar. 2019. [A theoretical analysis of contrastive unsupervised representation learning](#). In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California*, USA, volume 97 of *Proceedings of Machine Learning Research*, pages 5628–5637. PMLR.
- 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, Berlin, Germany. Association for Computational Linguistics.
- Darsh Shah, Tao Lei, Alessandro Moschitti, Salvatore Romeo, and Preslav Nakov. 2018. [Adversarial domain adaptation for duplicate question detection](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1056–1063, Brussels, Belgium. Association for Computational Linguistics.
- Jian Shen, Yanru Qu, Weinan Zhang, and Yong Yu. 2018. [Wasserstein distance guided representation learning for domain adaptation](#). In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 4058–4065. AAAI Press.
- Rui Wang, Zuxuan Wu, Zejia Weng, Jingjing Chen, Guo-Jun Qi, and Yu-Gang Jiang. 2021. [Cross-domain contrastive learning for unsupervised domain adaptation](#). *ArXiv preprint*, abs/2106.05528.
- Zirui Wang, Zihang Dai, Barnabás Póczos, and Jaime G. Carbonell. 2019. [Characterizing and avoiding negative transfer](#). In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 11293–11302. Computer Vision Foundation / IEEE.
- Xing Wu, Tao Zhang, Liangjun Zang, Jizhong Han, and Songlin Hu. 2019a. [Mask and infill: Applying masked language model for sentiment transfer](#). In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, pages 5271–5277. ijcai.org.
- Yifan Wu, Ezra Winston, Divyansh Kaushik, and Zachary C. Lipton. 2019b. [Domain adaptation with asymmetrically-relaxed distribution alignment](#). In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 6872–6881. PMLR.
- Zhuofeng Wu, Sinong Wang, Jiatao Gu, Madian Khabsa, Fei Sun, and Hao Ma. 2020. [Clear: Contrastive learning for sentence representation](#). *ArXiv preprint*, abs/2012.15466.

- Yuanmeng Yan, Rumei Li, Sirui Wang, Fuzheng Zhang, Wei Wu, and Weiran Xu. 2021. [ConSERT: A contrastive framework for self-supervised sentence representation transfer](#). 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 5065–5075, Online. Association for Computational Linguistics.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime G. Carbonell, Ruslan Salakhutdinov, and Quoc V. Le. 2019. [Xlnet: Generalized autoregressive pretraining for language understanding](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 5754–5764.
- David Yarowsky. 1995. [Unsupervised word sense disambiguation rivaling supervised methods](#). In *33rd Annual Meeting of the Association for Computational Linguistics*, pages 189–196, Cambridge, Massachusetts, USA. Association for Computational Linguistics.
- Hai Ye, Qingyu Tan, Ruidan He, Juntao Li, Hwee Tou Ng, and Lidong Bing. 2020. [Feature adaptation of pre-trained language models across languages and domains with robust self-training](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7386–7399, Online. Association for Computational Linguistics.
- Zhi-Hua Zhou and Ming Li. 2005. Tri-training: Exploiting unlabeled data using three classifiers. *IEEE Transactions on knowledge and Data Engineering*, 17(11):1529–1541.
- Chen Zhu, Yu Cheng, Zhe Gan, Siqi Sun, Tom Goldstein, and Jingjing Liu. 2020. [Freelb: Enhanced adversarial training for natural language understanding](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Fuzhen Zhuang, Xiaohu Cheng, Ping Luo, Sinno Jialin Pan, and Qing He. 2015. [Supervised representation learning: Transfer learning with deep autoencoders](#). In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, pages 4119–4125. AAAI Press.
- Yftah Ziser and Roi Reichart. 2017. [Neural structural correspondence learning for domain adaptation](#). In *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*, pages 400–410, Vancouver, Canada. Association for Computational Linguistics.

A Datasets

We obtain the Amazon review datasets from [He et al. \(2018\)](#). This dataset does not remove neutral labels and will not be problematic in UDA situation where the label information of the target domain is not available. In addition, reserving neutral labels also bring challenges for pre-trained language model, making it more favorable for self-supervised representation learning. Summary of this dataset is available in Table 4.

Domain		#Neg	#Neu	#Pos	Total
Book	Set 1	2000	2000	2000	6000
	Set 2	513	663	4824	6000
Electronics	Set 1	2000	2000	2000	6000
	Set 2	694	489	4817	6000
Beauty	Set 1	2000	2000	2000	6000
	Set 2	616	675	4709	6000
Music	Set 1	2000	2000	2000	6000
	Set 2	785	774	4441	6000

Table 4: Amazon review dataset

Each domain contains two sets, set 1 contains 6000 instances with balanced class labels, and set 2 contains instances that are randomly sampled from the larger dataset ([McAuley et al., 2015](#)), preserving authentic label distribution, examples in these two datasets do not overlap. Following ([He et al., 2018](#)), we use set 1 from the source domain as the training set for all our experiments. Since label distribution in the target domain is unpredictable and out of control in real life, so it’s more reasonable to use set 2 from the target domain as the unlabeled set, lastly the model will be evaluated in set 1 from target domain. For data split, we randomly sample 1000 instances from the source labeled dataset as validation set. When running UDA experiments, the model will train on 5000 source labeled examples and 6000 target unlabeled examples, then validate on 1000 source labeled examples.

For Amazon Benchmark([Blitzer et al., 2007](#)), it also contains four domains: Book (BK), DVD (D), Electronics (E) and Kitchen housewares (K). There are 2000 balanced labeled data for each domain, we randomly select 20000 unlabeled reviews for BK, D and E. For K, only 17856 unlabeled reviews are available, statistics of Amazon Benchmark can be find in Table.5. For data split, 1600 balanced samples are randomly sampled from the source labeled dataset, and 400 for validation.

Domains	Labeled	Unlabeled
Book	2000	973194
Dvd	2000	122438
Electronics	2000	21009
Kitchen & housewares	2000	17856

Table 5: Amazon Benchmark statistics

B DANN

Domain Adversarial Neural Network [Ganin et al. \(2016\)](#) proposes Domain Adversarial Neural Network (DANN), which learns domain invariant and discriminative features simultaneously. This approach is motivated by the idea that an adaptation algorithm could learn good representations for cross-domain transfer if it cannot differentiate the domain of the input observations. The optimization objective is:

$$\min_{\theta_f, \theta_y} \sum_{(x,y) \sim \mathcal{D}^S} [\mathcal{L}(f(x; \theta_f, \theta_y), y) + \lambda R_{\theta_f}], \quad (11)$$

$$R_{\theta_f} = \max_{\theta_d} \sum_{(x,d) \sim \mathcal{D}^S, \mathcal{D}^T} [-\mathcal{L}(f(x; \theta_f, \theta_d), d)], \quad (12)$$

where θ_d is the parameter corresponding to the computation of the domain classification, d is the domain label, R_{θ_f} is a regularizer weighted by λ . Objective (11) learns task discrimination by minimizing task classification loss and tries to make features similar across domains by maximizing the domain classification loss.

Although the domain classifier with parameters θ_d could perfectly classify different domains, the balance between two terms in (11) is hard to maintain. Hence, the training process becomes unstable, requiring an elaborate adaptation rate λ tuning ([Shen et al., 2018](#); [Shah et al., 2018](#); [Du et al., 2020](#)). Furthermore, the encoder could learn trivial solutions ([Karouzos et al., 2021](#)) which produce features with flipped domain predictions.