Improving Domain-Adapted Sentiment Classification by Deep Adversarial Mutual Learning

Qianming Xue,¹ Wei Zhang,^{1*} Hongyuan Zha²

¹School of Computer Science and Technology, KLATASDS-MOE, East China Normal University

²School of Computational Science and Engineering, Georgia Institute of Technology
xueqianming200@hotmail.com, zhangwei.thu2011@gmail.com, zha@cc.gatech.edu

Abstract

Domain-adapted sentiment classification refers to training on a labeled source domain to well infer document-level sentiment on an unlabeled target domain. Most existing relevant models involve a feature extractor and a sentiment classifier, where the feature extractor works towards learning domain-invariant features from both domains, and the sentiment classifier is trained only on the source domain to guide the feature extractor. As such, they lack a mechanism to use sentiment polarity lying in the target domain. To improve domain-adapted sentiment classification by learning sentiment from the target domain as well, we devise a novel deep adversarial mutual learning approach involving two groups of feature extractors, domain discriminators, sentiment classifiers, and label probers. The domain discriminators enable the feature extractors to obtain domain-invariant features. Meanwhile, the label prober in each group explores document sentiment polarity of the target domain through the sentiment prediction generated by the classifier in the peer group, and guides the learning of the feature extractor in its own group. The proposed approach achieves the mutual learning of the two groups in an end-to-end manner. Experiments on multiple public datasets indicate our method obtains the state-of-the-art performance, validating the effectiveness of mutual learning through label probers.

Introduction

Domain-adapted sentiment classification aims at training on a labeled source domain to well infer document-level sentiment on an unlabeled target domain. It is a natural intersection of researches on sentiment classification (Liu 2012) and unsupervised domain adaptation (Ben-David et al. 2010), and has attracted much attention with the flourish of online review platforms, such as Amazon, Yelp, etc. On the one hand, the total amount of reviews is exceedingly huge, which brings opportunities to pursue effective sentiment classification models (Tang, Qin, and Liu 2015a). On the other hand, a large number of review domains (e.g., product categories in Amazon) make it intractable to manually annotate enough data in each domain for training domain-specific models.

Thus developing automatically domain-adapted methods is imperative in this area. It is worth noting that this task is somewhat more challenging than some other cross-domain sentiment classification tasks which require a few labeled data in target domains (Peng et al. 2018).

A straightforward solution to the task is to directly apply the sentiment classification models trained on a source domain to a target domain. It does not obtain any training guidance from target domains. However, this is not ideal since it ignores the semantic gap between different domains. Taking Movie and Food domains as examples, the first domain contains common sentiment words such as “romantic”, “violent”, and “dramatic”. By contrast, the second includes opinion words like “tasty” and “delicious”. The potentially small set intersection of domain-dependent sentiment words indicates that there is huge potential for improving the naive solution. As such, it is important to build a domain-adapted sentiment classification model to perform well.

Existing relevant studies could be attributed into two categories: two-stage approaches (Blitzer, Dredze, and Pereira 2007; Glorot, Bordes, and Bengio 2011; Ziser and Reichart 2018; 2019) and end-to-end models (Ganin et al. 2016; He et al. 2018; Qu et al. 2019). The two-stage approaches typically construct unsupervised feature extractors or manually select pivot features across domains in the first stage. And in the second stage, a sentiment classifier is trained on the labeled source domain. Yet the first stage could not be directly guided by the ground-truth, and the selection of pivot features is a little too empirical and costly. On the other hand, benefiting from the advanced learning techniques such as adversarial learning (Ganin and Lempitsky 2015) and maximum mean discrepancy (Gretton et al. 2012a), end-to-end models overcome the above issues by training domain-variant feature extractors and sentiment classifiers holistically, without relying on pivot features. However, despite the promising results they have made, there still remains a major limitation in most of the end-to-end models: data from target domains is not fully utilized. That is, their sentiment classifiers and feature extractors overlook the sentiment polarity lying in the review text of target domains. One relevant study (He et al. 2018) in this regard obtains pseudo-labels in target domains by a self-ensemble bootstrapping technique

*Wei Zhang is corresponding author.

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

to train its sentiment classifier and feature extractor. However, the pseudo-labels are asynchronously generated by the earlier version of the sentiment classifier, which is a weaker classifier compared with its current version, and possibly limits the effectiveness of training.

In this paper, we propose a novel learning approach, named deep adversarial mutual learning (DAML). DAML improves domain-adapted sentiment classification by learning sentiment polarity from an unlabeled target domain. This is partially inspired by the recently proposed deep mutual learning (DML) (Zhang et al. 2018) for supervised single-domain tasks, where two classification models teach each other through their synchronously inferred sentiment label distributions which complement true labels. The rationality behind mutual learning is that these models can learn collaboratively and transfer their own learned knowledge to each other throughout the training process. This makes the models robust to the noise in the data. We extend mutual learning to the unsupervised cross-domain sentiment classification scenario through DAML. It involves two groups of feature extractors, domain discriminators, sentiment classifiers, and label probers. Different from standard mutual learning, we leverage the label prober in each group to learn pseudo sentiment label distributions of documents in the target domain, which are generated by the classifier from the other group. Meanwhile, the label probers guide the learning of feature extractors in their corresponding groups by gradient back-propagation. Through the above manner, probers act as bridges between classifiers and extractors, ensuring sentiment information in the target domain to be used by them. Another advantage of probers is that they free classifiers from aligning with each other, which is required by standard mutual learning and harms performance in a fully unsupervised scenario (see Figure 3a). In addition, we leverage gradient reverse layers (GRL) (Ganin et al. 2016) to learn domain-variant feature extractors with domain discriminators. The two groups can be regarded as two models which are mutually learned in an end-to-end manner to improve generalization on the target domain. We summarize the main contributions of this work as follows:

- We address the learning of sentiment polarity in the target domain, which is largely overlooked by previous studies. We propose DAML which combines the merits of adversarial learning and mutual learning, as well as the introduced label probers which learn from classifiers and guide the learning of extractors. To our knowledge, this is the first study of adapting mutual learning for domain-adapted sentiment classification.
- We evaluate DAML on multiple datasets with different origins. The extensive experiments show that DAML achieves the state-of-the-art performance. We also validate the better results of mutual learning with label probers than directly applying standard mutual learning. As a byproduct, we will release the source code of our approach¹.

¹<https://github.com/SleepyBag/DAML>

Related Work

In this section, we briefly discuss domain adaptation, sentiment classification, and deep mutual learning, to highlight the key differences in this research.

Domain adaptation. Domain adaptation has been a long-standing attractive research topic due to its real applications where labeled data is only available in a source domain (Pan and Yang 2009). It is widely recognized that the distribution gap between source domain and target domain is the fundamental challenge. To address this, early instance-based approaches reweigh each source example following the idea of importance sampling to match the target data distribution (Huang et al. 2007). Recent advancements in learning domain-variant feature representations include adversarial learning (Ganin and Lempitsky 2015) which trains an extractor to fool a domain classifier, and maximum mean discrepancy (MMD) (Gretton et al. 2012b) which measures the degree of domain shift and minimizes it.

In the specific scenario of domain-adapted sentiment classification, pivot-based methods (Blitzer, Dredze, and Pereira 2007; Yu and Jiang 2016; Ziser and Reichart 2018; 2019) first heuristically select domain-shared pivot features and then use them to learn the correspondence of domain-specific sentiment words. However, such an empirical selection is a little costly and its inaccuracy would be delivered to the following learning step. Some other studies such as (Glorot, Bordes, and Bengio 2011) adopt a similar two-stage procedure by learning an unsupervised domain-variant feature extractor (e.g., stacked denoising auto-encoder (Vincent et al. 2008)) in the first stage and then train a sentiment classifier by taking the obtained features as input. Due to the nature that the first stage of the above methods is not guided by sentiment labels, there is a potential to improve them by learning in an end-to-end manner. Some recent approaches (Ganin et al. 2016; He et al. 2018; Qu et al. 2019) employ adversarial learning or MMD to fulfill the end-to-end learning. However, as discussed previously, they inevitably suffer from the limitations of ignoring sentiment polarity lying in target domains or not leveraging them in an effective manner. In comparison, our work proposes the novel DAML approach to address the above issue, which is the most important contribution of this paper.

Sentiment classification. Despite the aforementioned studies about domain-adapted sentiment classification, a large amount of efforts have been devoted to the development of well-performed single-domain sentiment classification models (Zhang and Wang 2015), especially deep learning based ones, including recursive neural networks (Socher et al. 2013), convolutional neural networks (Kim 2014), and recurrent neural networks (Tang, Qin, and Liu 2015a), etc. To our knowledge, hierarchical attention networks (HAN) (Yang et al. 2016) is the state-of-the-art sentiment classification model, according to the performance comparison shown in (Chen et al. 2016; Wu et al. 2018). As such, our work takes the main part of HAN, except the last feed-forward output layer, as our feature extractor.

Deep mutual learning. Deep mutual learning (Zhang et al. 2018) is a recently developed learning approach in the background of model distillation (Hinton, Vinyals, and

Dean 2015), but it does not use a larger teacher model for guiding the training process. Apart from learning to fit the true labels, DML makes classifiers simultaneously learn from each other by mimicking others' inferred label distributions and finds its applications (Kanaci et al. 2019; Wu et al. 2019). However, none of existing studies have investigated its power in domain-adapted classification tasks. Later we show in the experiments that a naive combination of adversarial learning with standard mutual learning does not improve the performance in the task.

Our Approach

In domain-adapted sentiment classification, we are given a labeled source domain $\mathcal{D}_s = \{(\mathbf{x}_i^s, \mathbf{y}_i^s)\}_{i=1}^{n^s}$ with n^s documents and an unlabeled target domain $\mathcal{D}_t = \{\mathbf{x}_i^t\}_{i=n^s+1}^{n^s+n^t}$ with n^t documents, where $\mathbf{x}$ is the original word sequence representation in a document and $\mathbf{y}$ is the corresponding one-hot encoding of sentiment label. The goal of the task is to learn a well-performed sentiment classification model which is adapted to the target domain by leveraging the examples from $\mathcal{D}_s$ and $\mathcal{D}_t$. To this end, we present a deep adversarial mutual learning approach with the architecture shown in Figure 1.

Approach Overview: The proposed approach consists of two groups of feature extractors, domain discriminators, sentiment classifiers, and label probers. The two groups have exactly the same structures but are with their own parameters, wherein each group is associated with both the source and target domains. The feature extractors take documents from the source domain and the target domain as input, and eventually obtain document representations which are fed into the other components in each group. Domain discriminators, together with gradient reverse layers, ensure the extractors to get domain-invariant features by the guidance of domain discriminative loss $\mathcal{L}_{DOM}$. The prober in each group is associated with the classifier in the other group through the optimization of mutual learning based loss $\mathcal{L}_{ML}$. In addition, each classifier also corresponds to a classification loss $\mathcal{L}_{CLS}$ w.r.t. the labeled source domain.

Feature Extraction

We use HAN as the feature extractors to produce the representations of sentiment documents. HAN takes a hierarchical structure to mirror the word-level and sentence-level structures of documents. Assume the i -th document from either $\mathcal{D}_s$ or $\mathcal{D}_t$ contains l_i sentences and each sentence $j \in \{1, \dots, l_i\}$ has $m_{i,j}$ words. HAN first leverages a bidirectional GRU (Bahdanau, Cho, and Bengio 2015) to get contextualized word-level representations as follows,

$$\mathbf{h}_{i,j,k} = \text{BiGRU}(\mathbf{x}_{i,j,k}), k \in \{1, \dots, m_{i,j}\}. \quad (1)$$

To capture the importance of each word in constructing a sentence vector $\mathbf{s}_{i,j}$, word-level attention mechanism is used and defined as follows:

$$\alpha_{i,j,k} = \frac{\exp(\omega_w^\top \tanh(\mathbf{W}_w \mathbf{h}_{i,j,k} + \mathbf{b}_w))}{\sum_{k'} \exp(\omega_w^\top \tanh(\mathbf{W}_w \mathbf{h}_{i,j,k'} + \mathbf{b}_w))},$$

$$\mathbf{s}_{i,j} = \sum_k \alpha_{i,j,k} \mathbf{h}_{i,j,k}, \quad (2)$$

Figure 1: The architecture of deep adversarial mutual learning. The red lines and blues lines indicate the information flow from source domain and the target domain, respectively, while the black lines correspond to both domains. $\mathcal{L}_{DOM}$, $\mathcal{L}_{ML}$, and $\mathcal{L}_{CLS}$ denotes domain discriminative loss, mutual learning based loss, and source domain sentiment classification loss, respectively.

where $\alpha_{i,j,k}$ quantifies the word importance weight.

Afterwards, another bidirectional GRU is used to model $\mathbf{s}_{i,j}$ to get contextualized sentence representations $\mathbf{h}_{i,j}$ ($j \in \{1, \dots, l_i\}$). Given this, the document-level representation $\mathbf{d}_i$ is obtained through sentence-level attention mechanism in a similar fashion as Eq. 2, i.e., $\mathbf{d}_i = \sum_j \alpha_{i,j} \mathbf{h}_{i,j}$.

To summarize the procedure of feature extraction, we denote it as $\mathbf{d}_i = \text{FE}(\mathbf{x}_i; \Theta_{FE})$ where Θ_{FE} covers all the parameters involved in the above computations.

Sentiment Classification

We define a sentiment classifier $C(\mathbf{d}_i; \Theta_C)$ which contains one to several layers of perceptrons and is given as follows:

$$C(\mathbf{d}_i; \Theta_C) = \text{softmax}(\mathbf{W}_c \mathbf{d}_i + \mathbf{b}_c), \quad (3)$$

where $\mathbf{W}_c$ and $\mathbf{b}_c$ are learnable parameters and the dimension of output equals to the number of sentiment labels. Given this, the classification loss on the labeled source domain is based on cross-entropy, given as:

$$\mathcal{L}_{CLS} = \sum_{i=1}^{n^s} \mathbf{y}_i^s \log(C(\mathbf{d}_i)). \quad (4)$$

Domain Adaptation

To empower the feature extractors with the ability of learning domain-invariant representations, we consider the adversarial learning method. In particular, we introduce a domain discriminator D , which is trying to figure out from which domain a document vector comes. On the contrary, the feature extractor FE aims to fool D .

We define D as a multi-layer perceptron with parameters Θ_D , which regards the previously obtained document vector as input and outputs a scalar indicating the probability of a document being from the source domain. For document $\mathbf{x}_i$, we set $z_i = 1$ if it belongs to the source domain, and $z_i = 0$ for the target domain. Based on this, a min-max game is played to optimize the parameters Θ_{FE} and Θ_D as follows:

$$\Theta_{FE}, \Theta_D = \underset{\Theta_{FE}}{\operatorname{argmax}} \underset{\Theta_D}{\min} \mathcal{L}_{DOM} \quad (5)$$

where $\mathcal{L}_{DOM} = -\sum_{i=1}^{n^s+n^t} (z_i \log D(\mathbf{d}_i) + (1-z_i) \log(1-D(\mathbf{d}_i)))$.

As suggested in (Ganin et al. 2016), this min-max game is implemented by a gradient reversal layer, which reverses $\frac{\partial \mathcal{L}_{DOM}}{\partial \Theta_{FE}}$ into $-\eta \frac{\partial \mathcal{L}_{DOM}}{\partial \Theta_{FE}}$ during the gradient back-propagation process. η is a controllable hyper-parameter. Once FE is able to prevent the discriminator from distinguishing data from the source or target domain, it is supposed to be successfully trained to retain only domain-independent information to some extent.

Integration with Mutual Learning

So far, we are able to utilize the target domain from the aspect of modeling its data distribution. However, it is still far from satisfactory since the involved sentiment information is not explicitly captured. A natural solution is to generate pseudo-labels for the target domain to guide the learning of the sentiment classification model.

Motivated by mutual learning (Zhang et al. 2018) in supervised single-domain tasks, we first consider a direct way of incorporating standard mutual learning (sML) into our framework. We extend the aforementioned feature extractor, sentiment classifier, and domain discriminator to form two groups (models), i.e., $G_1 = (FE_1, C_1, D_1; \Theta_{G_1})$ and $G_2 = (FE_2, C_2, D_2; \Theta_{G_2})$, where Θ_{G_1} and Θ_{G_2} cover all the parameters of the groups, i.e., $\Theta_{G_1} = \{\Theta_{FE_1}, \Theta_{C_1}, \Theta_{D_1}\}$ and $\Theta_{G_2} = \{\Theta_{FE_2}, \Theta_{C_2}, \Theta_{D_2}\}$. Correspondingly, they have $\mathcal{L}_{CLS_1}$ and $\mathcal{L}_{DOM_1}$, $\mathcal{L}_{CLS_2}$ and $\mathcal{L}_{DOM_2}$, respectively. In consequence, we define the loss of standard mutual learning in our problem setting as follows:

$$\mathcal{L}_{sML_1} = \sum_{i=n^s+1}^{n^s+n^t} D_{KL}(C_2(\mathbf{d}_i) \| C_1(\mathbf{d}_i)), \quad (6)$$

$$\mathcal{L}_{sML_2} = \sum_{i=n^s+1}^{n^s+n^t} D_{KL}(C_1(\mathbf{d}_i) \| C_2(\mathbf{d}_i)), \quad (7)$$

where D_{KL} is Kullback Leibler (KL) divergence. Moreover, the two groups have their own hybrid objectives which are defined as follows:

$$\mathcal{L}_{G_1} = \mathcal{L}_{CLS_1} + \lambda_D \mathcal{L}_{DOM_1} + \lambda_M \mathcal{L}_{sML_1}, \quad (8)$$

$$\mathcal{L}_{G_2} = \mathcal{L}_{CLS_2} + \lambda_D \mathcal{L}_{DOM_2} + \lambda_M \mathcal{L}_{sML_2}, \quad (9)$$

where λ_D and λ_M are parameters to control the relative influence of the losses in the whole objectives. An alternating optimization strategy is used for mutually learning the two groups. That is, in each mini-batch, $\frac{\partial \mathcal{L}_{G_1}}{\partial \Theta_{G_1}}$ and $\frac{\partial \mathcal{L}_{G_2}}{\partial \Theta_{G_2}}$ are computed and propagated back to update model parameters.

It is intuitive that due to different parameter initialization, the feature extractor and sentiment classifier in each group try to obtain some different abilities (e.g., focusing on some special parts of latent feature spaces) to pursue similar results of its peer group, by the guidance of mutual learning. As such, they can transfer some exclusive knowledge to the peers. It is worth noting that Eq. 6 and 7 only guide each group with the other's predicted labels, but without telling them how to obtain these pseudo-labels, thus not forcing

them to obtain very similar feature representations. However, the way in which these pseudo-labels are applied to the unsupervised domain-adaptation scenario is somewhat questionable. This is because the requirement that the two classifiers are made to trust the pseudo-labels generated by each other (see Eq. 6 and 7) is a little too restricted in domain adaptation, where the two classifiers actually cannot see any labeled data in the target domain. As a result, the classifiers might be not strong enough to teach each other, which turns out that they are misled by the counterparts' predictions and cannot achieve improved classification performance (see Figure 3 for validation).

To avoid to degrade the classification performance and meanwhile retain the ability of learning sentiment polarity from the target domain, we introduce another component called label prober (P) to each group, which has the same architecture as the sentiment classifier but comes with its own learnable parameters. The goal of label probers is not for sentiment classification like sentiment classifiers. Instead, each prober probes sentiment information learned by the other group in the target domain, and transfers the knowledge to the feature extractor from the same group, acting as a bridge which connects the feature extractor to the sentiment classifier from the other group. In particularly, our mutual learning based loss functions are defined as follows:

$$\mathcal{L}_{ML_1} = \sum_{i=n^s+1}^{n^s+n^t} D_{KL}(C_2(\mathbf{d}_i) \| P_1(\mathbf{d}_i)), \quad (10)$$

$$\mathcal{L}_{ML_2} = \sum_{i=n^s+1}^{n^s+n^t} D_{KL}(C_1(\mathbf{d}_i) \| P_2(\mathbf{d}_i)). \quad (11)$$

Correspondingly, we have two new versions of objectives $\mathcal{L}_{G_1}$ and $\mathcal{L}_{G_2}$ by replacing $\mathcal{L}_{sML_1}$ and $\mathcal{L}_{sML_2}$ with $\mathcal{L}_{ML_1}$ and $\mathcal{L}_{ML_2}$. Taking Θ_{C_1} for illustration, compared with standard mutual learning, our approach changes the gradients of the classifiers from $\frac{\partial \mathcal{L}_{CLS_1}}{\partial \Theta_{C_1}} + \lambda_M \frac{\partial \mathcal{L}_{sML_1}}{\partial \Theta_{C_1}}$ to $\frac{\partial \mathcal{L}_{CLS_1}}{\partial \Theta_{C_1}}$. Thus it frees the classifiers from learning not very reliable pseudo-labels as mentioned before.

Thanks to the integration of adversarial learning and mutual learning, the gradients of feature extractors are given as: $\frac{\partial \mathcal{L}_{G_1}}{\partial \Theta_{FE_1}} = \frac{\partial \mathcal{L}_{CLS_1}}{\partial \Theta_{FE_1}} - \eta \lambda_D \frac{\partial \mathcal{L}_{DOM_1}}{\partial \Theta_{FE_1}} + \lambda_M \frac{\partial \mathcal{L}_{ML_1}}{\partial \Theta_{FE_1}}$ and $\frac{\partial \mathcal{L}_{G_2}}{\partial \Theta_{FE_2}} = \frac{\partial \mathcal{L}_{CLS_2}}{\partial \Theta_{FE_2}} - \eta \lambda_D \frac{\partial \mathcal{L}_{DOM_2}}{\partial \Theta_{FE_2}} + \lambda_M \frac{\partial \mathcal{L}_{ML_2}}{\partial \Theta_{FE_2}}$, enabling the feature extractors to learn data distributions and sentiment information from both the source domain and the target domain.

Experiments

In this section, we assess the effectiveness of our approach DAML by first clarifying the experimental setup and afterwards analyzing the experimental results.

Experimental Setup

Datasets We adopt multiple publicly available datasets with different origins to evaluate DAML. The first pair of source and target datasets is Yelp and IMDB datasets built

Domain	Training	Development	Test
Yelp	62,522	7,773	8,671
IMDB	67,426	8,381	9,112
Electronics	79,942	9,994	9,994
CD	79,999	10,000	10,000
Clothing	79,990	10,000	10,000

Table 1: Statistics of the experimental datasets.

by (Tang, Qin, and Liu 2015b). One issue is that Yelp has 5 sentiment labels, while the number is 10 for IMDB. To align the space of sentiment labels for domain adaptation, we simply divide the scores of the sentiment labels in IMDB by 2 and round them up.

Moreover, to investigate the performance in different domains with the same origin, we choose three domains, i.e., Electronics, CD, and Clothing, from the Amazon dataset (McAuley et al. 2015). For each domain, there are 80,000 pieces of documents in the training set and 10,000 in both the test set and development set. All of these statistics are summarized in Table 1 in detail.

Evaluation Following the study of (Tang, Qin, and Liu 2015b), the model performance in our experiments is compared according to *Acc*, short for *Accuracy*, which measures the classification performance, and *RMSE*, which shows the divergence between the predicted ratings and the ground truth ratings.

For unsupervised domain-adaptation task, the unavailability of the development set in target domains results in two issues. First, it is impossible to tune hyper-parameters for every task according to the performance on the development set of each target domain. Therefore, as is done by (He et al. 2018), we build a development set available for only one task (Yelp $\rightarrow$ IMDB). Hyper-parameters of our model and baselines are tuned to perform their best on this development set and then fixed on all tasks. Second, we are not able to do early-stopping for a model according to its performance in a target domain. As an alternative, we use the development set of each source domain to determine early-stopping. Specifically, for each model, we save its parameters when it achieves better results on the development set along the training process. And finally the saved parameters with best performance are regarded as model parameters for testing on the target domain. All the models experimented in this work determine their model parameters through the above manner.

For each mutual learning based approach, the classifier that performs the best on the development set of a source domain is chosen to be eventually evaluated on a target domain.

Implementations As mentioned in the section of feature extraction, our approach uses HAN as the feature extractor for its good performance. For a fair comparison, we make other approaches to use HAN as well. Though Bidirectional Encoder Representations from Transformers (BERT) (Devlin et al. 2019) has been gaining more and more attention as a powerful text feature extractor, we did not find an appro-

priate setting to obtain significantly better results than HAN from our local experiments.

Before the training process starts, 200-dimensional word embeddings are learned by Word2vec (Mikolov et al. 2013). Word2vec is trained on all the available data in all the experimented domains. The hyper-parameters are fixed for all domains. By default, η , λ_D , λ_M are set as 0.005, 1.0 and 1.0, respectively. Adam is adopted as the optimizer for all the experiments.

For the baselines, if their source codes are available, we directly utilize them by just adjusting input and tuning parameters. Otherwise, we implement them by referring to the settings shown in their papers.

Baselines The following end-to-end domain-adapted sentiment classification models are adopted for comparison.

Naive: We experiment with a sentiment classification model implemented by HAN without any domain-adaptation technique. It is trained on a source domain and then directly tested on a target domain.

DANN (Ganin et al. 2016): An adversarial discriminator is introduced to distinguish data from different domains so that its feature extractor could finally generate domain-invariant latent feature vectors by trying to confuse the discriminator.

MMD (Gretton et al. 2012a): As another common method for domain adaptation, MMD is applied to a feature extractor to make it domain-independent. The results of MMD are mainly compared with the ones of adversarial domain adaptation and so as to examine the effectiveness of the latter.

WDGRL (Shen et al. 2017): This approach is similar to DANN, but uses Wasserstein Distance instead of JS-divergence as the loss function of the discriminator.

DAS (He et al. 2018)²: This approach employs entropy minimization and a self-ensemble bootstrapping method to refine its classifier while minimizing the domain divergence.

ACAN (Qu et al. 2019)³: ACAN introduces two classifier networks on the top of the feature extractor. The discrepancy of two classifiers is increased to provide diverse views, while the extractor learns to create better features away from the category boundaries by minimizing this discrepancy.

To fully investigate the mechanism behind our approach, we consider the following variants:

Naive Ensemble (NE): To show that DAML performs well in terms of leveraging multiple parallel models, we design a naive ensemble framework which fetches the probabilities returned by two domain-adapted models and takes their addition as its prediction.

ML (Zhang et al. 2018): It is the direct application of standard mutual learning with adversarial learning to our task, as discussed in previous section.

Feature Alignment (FA): This method aligns two classification models at the level of latent feature layer by minimizing the Euclidean distance between the outputs of two feature extractors for each input document.

²<https://github.com/ruidan/DAS>

³<https://github.com/XiaoYee/ACAN>

Table 2: Sentiment classification results in terms of Acc and RMSE. The best results are marked in bold.

Source ↔ Target	Metrics	Baselines						Our Approaches			
		Naive	DANN	MMD	WDGRL	DAS	ACAN	NE	ML	FA	DAML
Yelp	Acc	0.513	0.528	0.525	0.522	0.471	0.508	0.532	0.529	0.519	0.563
↔ IMDB	RMSE	0.956	0.867	0.877	0.882	1.029	0.967	0.867	0.907	0.925	0.822
IMDB	Acc	0.510	0.522	0.511	0.494	0.509	0.502	0.518	0.517	0.509	0.528
↔ Yelp	RMSE	0.921	0.881	0.971	0.955	0.929	0.964	0.865	0.889	0.883	0.904
CD	Acc	0.662	0.664	0.655	0.655	0.627	0.657	0.653	0.660	0.654	0.669
↔ Electronics	RMSE	1.074	1.016	1.046	1.097	1.123	1.109	1.077	1.030	1.059	1.068
Electronics	Acc	0.654	0.654	0.649	0.655	0.628	0.654	0.659	0.657	0.655	0.670
↔ CD	RMSE	1.020	0.970	1.021	0.986	1.086	1.023	0.943	0.960	1.003	0.922
CD	Acc	0.658	0.657	0.651	0.648	0.582	0.656	0.660	0.622	0.621	0.667
↔ Clothing	RMSE	0.920	0.975	1.044	1.009	1.359	0.974	0.993	1.190	1.180	0.953
Clothing	Acc	0.644	0.639	0.650	0.634	0.596	0.650	0.650	0.610	0.610	0.661
↔ CD	RMSE	1.067	1.034	1.021	0.919	1.285	1.030	0.957	1.246	1.229	0.917
Electronics	Acc	0.692	0.693	0.689	0.693	0.582	0.693	0.702	0.690	0.672	0.705
↔ Clothing	RMSE	0.819	0.851	0.840	0.805	1.359	0.827	0.806	0.857	0.929	0.781
Clothing	Acc	0.675	0.689	0.680	0.678	0.594	0.685	0.696	0.663	0.666	0.698
↔ Electronics	RMSE	0.878	0.883	0.911	0.913	1.424	0.887	0.891	1.086	1.033	0.819
Average	Acc	0.626	0.631	0.626	0.622	0.547	0.626	0.634	0.619	0.613	0.645
Performance	RMSE	0.957	0.935	0.966	0.946	1.199	0.972	0.925	1.027	1.030	0.898

Experimental Results

Model comparison Table 2 compares our approaches with baselines on different domain-adaptation tasks. Our final approach DAML outperforms all the other approaches on most tasks. Compared with the NE framework, which in average performs the best among all the left approaches, DAML improves from 63.4% to 64.5% in terms of accuracy and reduces the RMSE error from 0.925 to 0.898. Moreover, NE requires multiple available classification models when testing, while DAML only needs a single model that performs the best on the development set of a source domain.

Among the four proposed methods, the ML and FA learning frameworks perform significantly worse than NE and DAML, and even worse than single models. As is analyzed before, ML requires each model to trust its weak peer. Its failure emphasizes the intuition that models ought to take different views on ground truth and pseudo-labels. On the other hand, though FA tries to align multiple models on the layer of the feature representation without disturbing their classifiers, it does not provide models with capability on learning from each other effectively.

Mutual learning on different domain settings So far, one of the most important remaining questions is whether our learning framework works by offering pseudo-labels of the target domain or by only introducing a regularizer that pulls two models closer. To answer this question, we try to train a group of the DAML frameworks which apply probers to only the source domain, only the target domain, and both the two domains. These experiments are conducted on the Yelp, IMDB, Amazon Electronics, and CD datasets.

Table 3 presents the corresponding results. We can see training the probers of DAML on the source domain exclusively performs much worse than the ones trained on the target domain, or both domains. The comparison emphasizes

Table 3: Results of DAML trained on different domains, with its probers exclusively applied to the domain indicated by the column name.

Source ↔ Target	Metrics	Target Domain	Source Domain	Both Domains
Yelp	Acc	0.563	0.521	0.546
↔ IMDB	RMSE	0.822	0.895	0.827
IMDB	Acc	0.528	0.514	0.531
↔ Yelp	RMSE	0.904	0.890	0.852
CD	Acc	0.669	0.652	0.663
↔ Electronics	RMSE	1.068	1.085	1.015
Electronics	Acc	0.670	0.669	0.664
↔ CD	RMSE	0.922	0.942	0.946
Average	Acc	0.608	0.589	0.601
Performance	RMSE	0.929	0.953	0.910

the importance of leveraging the sentiment information in the target domains. According to the results, we can conclude that what makes DAML perform well is its ability of taking an advantage of pseudo-labels.

Feature visualization Taking the task of Yelp → IMDB for illustration, we visualize the feature vectors of the target domain generated by our approach and standard mutual learning (ML). As shown in Figure 2, the feature representations learned by ML have vague classification boundaries. By contrast, our approach successfully reveals the sentiment polarity in the obtained representations, since negative (red and blue), neutral (green), and positive (purple and yellow) examples are located in the left, middle, and right parts of the figure, respectively.

Analysis on training process It can help understand the effect of each mutual learning framework to see its per-

(a) Target feature: Standard ML. (b) Target feature: DAML.

Figure 2: t-SNE visualization of the features in target domain learned by standard Mutual Learning and proposed DAML methods. Different rating scores are indicated by colors and shapes. (Red = 1, Blue = 2, Green = 3, Purple = 4, Yellow = 5)

(a) Standard ML vs. DANN. (b) DAML vs. DANN.

Figure 3: Accuracy on the test set, evaluated at each training step. In each picture, two solid lines stand for the two models trained in the framework. The dashed line represents the plain adversarial domain-adaptation method DANN.

formance in the target domain at each step. Therefore, we conduct a series of experiments on the Yelp→IMDB task, where a development set from the target domain is available. The performance on the development set of the chosen approaches is evaluated every several steps and the curves of accuracy are drawn in Figure 3.

We can see DANN and two sentiment classifiers in the two mutual learning based frameworks obtain high performance at about the twentieth steps. After that, DANN keeps its performance relatively stable, compared with two other models in the figure. Particularly, the curves of standard ML start to drop rapidly, while our final approach gradually obtains even higher performance. These comparisons clearly illustrate the negative influence of trusting each other’s weak domain-adapted classifier to much.

Influence of hyper-parameter η and λ_M In order to observe the influence of the relative weights of different parts in the total loss function, we conduct some experiments with the development sets of the target domains. Specifically, three tasks, i.e., Yelp → IMDB, Clothing → Electronics and Electronics → CD, are chosen. For each task, our learning approach is trained and then validated on the development set of the target domain for just showing the influence. This procedure is repeated several times with different η and λ_M . λ_D is not considered because it only influences the performance of the discriminators when η controls the gradient

(a) Influence of η . (b) Influence of λ .

Figure 4: DAML trained on three tasks with different η and λ_M . The vertical axis stands for the relative accuracy value gained w.r.t. the minimum point for each curve.

flowed to the feature extractors.

Figure 4 depicts how the change of η and λ_M influences the performance. It can be seen from the left subfigure that DAML does not perform very well when adversarial learning is not applied (when η is 0), which shows the contribution of adversarial learning to our learning framework. Moreover, we can observe from the right subfigure that DAML obtains good performance when λ_M is in the range $[1.0, 5.0]$ for all the three tasks. Besides, compared with the standard domain-adaptation model without mutual learning (when λ_M equals to 0), DAML gains improvement when λ_M is in even a larger range. These curves show the effectiveness of our learning framework and its robustness in terms of the critical hyper-parameters.

Influence of number of groups and models In the end, we concisely analyze how the results change by increasing the number of groups in DAML and domain-adapted models in the baseline NE, from 2 to 4. We take the domains of Yelp and IMDB for testing and show the performance in Table 4.

Table 4: Sentiment classification results of more groups and models.

Source ↪ Target	Metrics	DAML			NE		
		2	3	4	2	3	4
Yelp ↪ IMDB	Acc	0.563	0.563	0.547	0.532	0.527	0.534
	RMSE	0.822	0.819	0.840	0.867	0.874	0.887
IMDB ↪ Yelp	Acc	0.528	0.531	0.522	0.518	0.515	0.506
	RMSE	0.904	0.874	0.916	0.865	0.903	0.876

It can be seen that, in this work, involving more models doesn’t always improve the performance of the approaches. This might be attributed to the fact that a development set from a target domain is hardly available in practice, as mentioned in the section of evaluation. Therefore, we follow previous work to train each model until it performs the best on a source domain. In consequence, though an approach with more groups or models may perform better on the target domain at some training points during the learning process, we are not able to finally obtain these peak points.

Conclusion

In this work, we have devised DAML, a novel deep adversarial mutual learning framework through which two domain-

adapted models can learn from each other to effectively utilize the sentiment information lying in a target domain. This advantage is enjoyed by the means of incorporating probes into mutual learning, which act as bridges to connect feature extractors with sentiment classifiers. Unlike standard mutual learning, our learning framework introduces a flexible manner for each model to decide how to view its peer's predictions, rather than to align the output labels of sentiment classifiers, which is somewhat questionable in unsupervised domain-adaptation tasks. Experiments have shown that DAML significantly improves the performance over representative end-to-end models for domain-adapted sentiment classification.

Acknowledgments

We thank the anonymous reviewers for their valuable comments. This work is supported by National Key Research and Development Program (2019YFB2102600), NSFC (61702190), NSFC-Zhejiang (U1609220), and Shanghai Sailing Program (17YF1404500).

References

- Bahdanau, D.; Cho, K.; and Bengio, Y. 2015. Neural machine translation by jointly learning to align and translate. In *ICLR*.
- Ben-David, S.; Blitzer, J.; Crammer, K.; Kulesza, A.; Pereira, F.; and Vaughan, J. W. 2010. A theory of learning from different domains. *Machine Learning* 79(1-2):151–175.
- Blitzer, J.; Dredze, M.; and Pereira, F. 2007. Biographies, bollywood, boom-boxes and blenders: Domain adaptation for sentiment classification. In *ACL*.
- Chen, H.; Sun, M.; Tu, C.; Lin, Y.; and Liu, Z. 2016. Neural sentiment classification with user and product attention. In *EMNLP*, 1650–1659.
- Devlin, J.; Chang, M.; Lee, K.; and Toutanova, K. 2019. BERT: pre-training of deep bidirectional transformers for language understanding. In *NAACL*, 4171–4186.
- Ganin, Y., and Lempitsky, V. S. 2015. Unsupervised domain adaptation by backpropagation. In *ICML*, 1180–1189.
- Ganin, Y.; Ustinova, E.; Ajakan, H.; Germain, P.; Larochelle, H.; Laviolette, F.; Marchand, M.; and Lempitsky, V. 2016. Domain-adversarial training of neural networks. *JMLR* 17(1):2096–2030.
- Glorot, X.; Bordes, A.; and Bengio, Y. 2011. Domain adaptation for large-scale sentiment classification: A deep learning approach. In *ICML*, 513–520.
- Gretton, A.; Borgwardt, K. M.; Rasch, M. J.; Schölkopf, B.; and Smola, A. 2012a. A kernel two-sample test. *JMLR* 13(Mar):723–773.
- Gretton, A.; Sriperumbudur, B. K.; Sejdinovic, D.; Strathmann, H.; Balakrishnan, S.; Pontil, M.; and Fukumizu, K. 2012b. Optimal kernel choice for large-scale two-sample tests. In *NIPS*, 1214–1222.
- He, R.; Lee, W. S.; Ng, H. T.; and Dahlmeier, D. 2018. Adaptive semi-supervised learning for cross-domain sentiment classification. In *EMNLP*, 3467–3476.
- Hinton, G.; Vinyals, O.; and Dean, J. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Huang, J.; Gretton, A.; Borgwardt, K.; Schölkopf, B.; and Smola, A. J. 2007. Correcting sample selection bias by unlabeled data. In *NIPS*, 601–608.
- Kanaci, A.; Li, M.; Gong, S.; and Rajamanoharan, G. 2019. Multi-task mutual learning for vehicle re-identification. In *CVPR Workshops*, 62–70.
- Kim, Y. 2014. Convolutional neural networks for sentence classification. In *EMNLP*, 1746–1751.
- Liu, B. 2012. *Sentiment Analysis and Opinion Mining*. Synthesis Lectures on Human Language Technologies. Morgan & Claypool Publishers.
- McAuley, J. J.; Targett, C.; Shi, Q.; and van den Hengel, A. 2015. Image-based recommendations on styles and substitutes. In *SIGIR*, 43–52.
- Mikolov, T.; Sutskever, I.; Chen, K.; Corrado, G. S.; and Dean, J. 2013. Distributed representations of words and phrases and their compositionality. In *NIPS*, 3111–3119.
- Pan, S. J., and Yang, Q. 2009. A survey on transfer learning. *IEEE TKDE* 22(10):1345–1359.
- Peng, M.; Zhang, Q.; Jiang, Y.; and Huang, X. 2018. Cross-domain sentiment classification with target domain specific information. In *ACL*, 2505–2513.
- Qu, X.; Zou, Z.; Cheng, Y.; Yang, Y.; and Zhou, P. 2019. Adversarial category alignment network for cross-domain sentiment classification. In *ACL*, 2496–2508.
- Shen, J.; Qu, Y.; Zhang, W.; and Yu, Y. 2017. Wasserstein distance guided representation learning for domain adaptation. *arXiv preprint arXiv:1707.01217*.
- Socher, R.; Perelygin, A.; Wu, J.; Chuang, J.; Manning, C. D.; Ng, A. Y.; and Potts, C. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *ACL*, 1631–1642.
- Tang, D.; Qin, B.; and Liu, T. 2015a. Document modeling with gated recurrent neural network for sentiment classification. In *EMNLP*, 1422–1432.
- Tang, D.; Qin, B.; and Liu, T. 2015b. Learning semantic representations of users and products for document level sentiment classification. In *ACL*, volume 1, 1014–1023.
- Vincent, P.; Larochelle, H.; Bengio, Y.; and Manzagol, P. 2008. Extracting and composing robust features with denoising autoencoders. In *ICML*, 1096–1103.
- Wu, Z.; Dai, X.; Yin, C.; Huang, S.; and Chen, J. 2018. Improving review representations with user attention and product attention for sentiment classification. In *AAAI*, 5989–5996.
- Wu, R.; Feng, M.; Guan, W.; Wang, D.; Lu, H.; and Ding, E. 2019. A mutual learning method for salient object detection with intertwined multi-supervision. In *CVPR*, 8150–8159.
- Yang, Z.; Yang, D.; Dyer, C.; He, X.; Smola, A.; and Hovy, E. 2016. Hierarchical attention networks for document classification. In *NAACL*, 1480–1489.
- Yu, J., and Jiang, J. 2016. Learning sentence embeddings with auxiliary tasks for cross-domain sentiment classification. In *EMNLP*, 236–246.
- Zhang, W., and Wang, J. 2015. Prior-based dual additive latent dirichlet allocation for user-item connected documents. In *IJCAI*, 1405–1411.
- Zhang, Y.; Xiang, T.; Hospedales, T. M.; and Lu, H. 2018. Deep mutual learning. In *CVPR*, 4320–4328.
- Ziser, Y., and Reichart, R. 2018. Pivot based language modeling for improved neural domain adaptation. In *NAACL*, 1241–1251.
- Ziser, Y., and Reichart, R. 2019. Task refinement learning for improved accuracy and stability of unsupervised domain adaptation. In *ACL*, 5895–5906.